

Communication-Efficient Algorithms for Statistical Optimization

Yuchen Zhang^{*†}John C. Duchi^{*†}Martin J. Wainwright^{*†‡}

September 20, 2012

Abstract

We present and study two communication-efficient algorithms for distributed statistical optimization on large-scale data sets. The first algorithm is a standard averaging method that distributes the N data samples evenly to m machines, performs separate minimization on each subset, and then averages the estimates. We provide a sharp analysis of this average mixture algorithm, showing that under a reasonable set of conditions, the combined parameter achieves mean-squared error that decays as $\mathcal{O}(N^{-1} + (N/m)^{-2})$. Whenever $m \leq \sqrt{N}$, this guarantee matches the best possible rate achievable by a centralized algorithm having access to all N samples. The second algorithm is a novel method, based on an appropriate form of bootstrap subsampling. Requiring only a single round of communication, it has mean-squared error that decays as $\mathcal{O}(N^{-1} + (N/m)^{-3})$, and so is more robust to the amount of parallelization. We also provide experimental evaluation of our methods, investigating their performance both on simulated data and on a large-scale regression problem from the internet search domain. In particular, we show that our methods can be used to efficiently solve an advertisement prediction problem from the Chinese SoSo Search Engine, which involves logistic regression with $N \approx 2.4 \times 10^8$ samples and $d \approx 740,000$ covariates.

1 Introduction

Many procedures for statistical estimation are based on a form of (regularized) empirical risk minimization, meaning that a parameter of interest is estimated by minimizing an objective function defined by the average of a loss function over the data. Given the current explosion in the size and amount of data available in statistical studies, a central challenge is to design efficient algorithms for solving large-scale problem instances. In a centralized setting, there are many procedures for solving empirical risk minimization problems, among them standard convex programming approaches [3] as well as stochastic approximation and optimization algorithms [27, 11, 20]. When the size of the dataset becomes extremely large, however, it may be infeasible to store all of the data on a single computer, or at least to keep the data in memory. Accordingly, the focus of this paper is the study of some distributed and communication-efficient procedures for empirical risk minimization.

^{*}Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

[†]Department of Statistics, University of California, Berkeley, Berkeley, CA 94720. Email: {yuczhang, jduchi, wainwrig}@eecs.berkeley.edu

[‡]YZ, JCD, and MJW were partially supported by MURI grant N00014-11-1-0688 from the Office of Naval Research. JCD was additionally supported by an NDSEG fellowship and a Facebook Fellowship.

Recent years have witnessed a flurry of research on distributed approaches to solving very large-scale statistical optimization problems. Although we cannot survey the literature adequately—the list of papers [19, 25, 12, 7, 6, 1, 26, 8] and references therein contain a sample of relevant work—we touch on a few important themes here. It can be difficult within a purely optimization-theoretic setting to show explicit benefits arising from distributed computation. In statistical settings, however, distributed computation can lead to gains in statistical efficiency, as shown by a number of authors [1, 6, 26, 8]. Within the family of distributed algorithms, there can be significant differences in communication complexity: different computers must be synchronized, and when the dimensionality of the data is high, communication can be prohibitively expensive. It is thus interesting to study distributed estimation algorithms that require fairly limited synchronization and communication, while still enjoying the greater statistical accuracy that is usually associated with a larger dataset.

With this context, perhaps the simplest algorithm for distributed statistical estimation is what we term the *average mixture* (AVGM) algorithm. It is an appealingly simple method: given m different machines and a dataset of size N , assign to each machine a (distinct) dataset of size $n = N/m$, have each machine i compute the empirical minimizer θ_i on its fraction of the data, and then average all the parameter estimates θ_i across the machines. This approach has been studied for some classification and estimation problems by Mann et al. [16] and McDonald et al. [18], as well as for certain stochastic approximation methods by Zinkevich et al. [30]. Given an empirical risk minimization algorithm that works on one machine, the procedure is straightforward to implement and is extremely communication efficient, requiring only a single round of communication. It is also relatively robust to possible failures in a subset of machines and/or differences in speeds, since there is no repeated synchronization. To the best of our knowledge, however, no work has shown rigorously that the AVGM procedure has greater statistical efficiency than the naive approach of using $n = N/m$ samples on a single machine.

This paper makes three main contributions. First, in Section 3 we provide a sharp analysis of the AVGM algorithm, showing that under a reasonable set of conditions on the statistical risk functional, it can indeed achieve substantially better rates than the naive approach. More concretely, we provide bounds on the mean-squared error that decay as $\mathcal{O}((nm)^{-1} + n^{-2})$. Whenever the number of machines m is less than the number of samples n per machine, this guarantee matches the best possible rate achievable by a centralized algorithm having access to all $N = nm$ samples. In the special case of optimizing log likelihoods, the pre-factor in our bound involves the trace of the Fisher information, a quantity well-known to control the fundamental limits of statistical estimation. We also show how the result extends to certain stochastic programming approaches, but we defer proofs of all results to later sections and appendices.

Our second contribution is to develop a novel extension of simple averaging; it is based on an appropriate form of bootstrap [9, 10, 22], which we refer to as the *bootstrap average mixture* (BAVGM) approach. At a high level, the BAVGM algorithm distributes samples evenly among m processors or computers as before, but instead of simply returning the empirical minimizer, each processor further subsamples its own dataset in order to estimate the bias of its own estimate, and returns a bootstrap-corrected estimate. We then prove that the BAVGM algorithm has mean-squared error decaying as $\mathcal{O}(m^{-1}n^{-1} + n^{-3})$. As long as $m < n^2$, the bootstrap method again matches the centralized gold standard up to higher order terms. These rates are sharper than any previous work, and moreover are minimax optimal up to constant factors.

Our third contribution is to perform a detailed empirical evaluation of both the AVGM and

BAVGM procedures, which we present in Sections 4 and 5. Using simulated data from normal and non-normal regression models, we explore the conditions under which the BAVGM algorithm yields better performance than the AVGM algorithm; in addition, we study the performance of both methods relative to an oracle baseline that uses all N samples. We also study the sensitivity of the algorithms to the number of splits m of the data, and in the BAVGM case, we investigate the sensitivity of the method to the amount of resampling. These simulations show that both AVGM and BAVGM have favorable performance, even when compared to the unattainable “gold standard” procedure that has access to all N samples. In Section 5, we complement our simulation experiments with a large logistic regression experiment that arises from the problem of predicting whether a user of a search engine will click on an advertisement. This experiment is large enough—involving $N \approx 2.4 \times 10^8$ samples in $d \approx 740,000$ dimensions with a storage size of approximately 55 gigabytes—that it cannot be solved efficiently on one machine. Consequently, a distributed approach is essential to take full advantage of this data set. Our experiments on this real-world problem show that the resampling and correction that the BAVGM method provides substantial performance benefits over naive solutions as well as the averaging algorithm AVGM. We provide our conclusions and discussion in Section 6.

2 Problem set-up

In this section, we formally describe our setting and the statistical optimization problems we consider, providing necessary background and defining notation as needed. We begin by setting up our decision-theoretic framework for empirical risk minimization, after which we describe our algorithms and the assumptions we require for our main theoretical results.

2.1 Empirical risk minimization

Let $\{f(\cdot; x), x \in \mathcal{X}\}$ be a collection of real-valued and convex loss functions with domain containing the convex set $\Theta \subseteq \mathbb{R}^d$. Let P be a probability distribution over the sample space \mathcal{X} , and define the *population risk functional* $F_0 : \Theta \rightarrow \mathbb{R}$

$$F_0(\theta) := \mathbb{E}_P[f(\theta; X)] = \int_{\mathcal{X}} f(\theta; x) dP(x). \quad (1)$$

Our goal is to estimate the parameter vector minimizing the population risk, namely the quantity

$$\theta^* := \operatorname{argmin}_{\theta \in \Theta} F_0(\theta) = \int_{\mathcal{X}} f(\theta; x) dP(x), \quad (2)$$

which we assume to be unique. In practice, the population distribution P is unknown to us, but we have access to a collection S of samples from the distribution P . Empirical risk minimization is based on estimating θ^* by solving the optimization problem

$$\hat{\theta} \in \operatorname{argmin}_{\theta \in \Theta} \left\{ \frac{1}{|S|} \sum_{x \in S} f(\theta; x) \right\}. \quad (3)$$

Throughout the paper, we impose some standard regularity conditions on the parameter space, the risk function F_0 , and the instantaneous loss functions $f(\cdot; x) : \Theta \rightarrow \mathbb{R}$. These conditions are natural and standard in classical statistical analysis of M -estimators (e.g., [15, 13]). Our first assumption deals with the relationship of the parameter space to the optimal parameter θ^* .

Assumption A (Parameters). *The parameter space $\Theta \subset \mathbb{R}^d$ is a compact convex set, with $\theta^* \in \text{int } \Theta$ and ℓ_2 -radius $R = \max_{\theta \in \Theta} \|\theta - \theta^*\|_2$.*

In addition, the risk function is required to have some amount of curvature. We formalize this notion as follows:

Assumption B (Local strong convexity). *There exists some $\lambda > 0$ such that the Hessian matrix of F_0 satisfies $\nabla^2 F_0(\theta^*) \succeq \lambda I_{d \times d}$.*

Here $\nabla^2 F_0(\theta)$ denotes the $d \times d$ Hessian matrix of the population objective F_0 evaluated at θ , and we use \succeq to denote the positive semidefinite ordering (i.e., $A \succeq B$ means that $A - B$ is positive semidefinite.) This local condition is milder than a global strong convexity condition and is required to hold only for the population risk F_0 evaluated at θ^* . Of course, some type of curvature of the risk is required to consistently estimate the parameters θ^* . (Without some local curvature, there may be no unique optimal parameter θ^* , so the model would be—in a sense—ill-specified.)

2.2 Averaging methods

In the distributed setting, we are given a dataset of $N = mn$ samples, drawn i.i.d. according to the initial distribution P , which we divide evenly and randomly among a total of m processors. (For simplicity, we assume the total number of samples is a multiple of m .) For $i = 1, \dots, m$, we let $S_{1,i}$ denote the data set assigned to processor i ; by construction, it is a collection of n samples drawn i.i.d. according to P , and the samples in subsets $S_{1,i}$ are $S_{1,j}$ are independent for $i \neq j$. In addition, for each processor i we define the (local) empirical distribution $P_{1,i}$ and empirical objective $F_{1,i}$ via

$$P_{1,i} := \frac{1}{|S_{1,i}|} \sum_{x \in S_{1,i}} \delta_x, \quad \text{and} \quad F_{1,i}(\theta) := \frac{1}{|S_{1,i}|} \sum_{x \in S_{1,i}} f(\theta; x). \quad (4)$$

With this notation, the AVGM algorithm is very simple to describe:

Average mixture algorithm

- (1) For each $i \in \{1, \dots, m\}$, processor i uses its local dataset $S_{1,i}$ to compute the local empirical minimizer

$$\theta_{1,i} \in \underset{\theta \in \Theta}{\operatorname{argmin}} \left\{ \underbrace{\frac{1}{|S_{1,i}|} \sum_{x \in S_{1,i}} f(\theta; x)}_{F_{1,i}(\theta)} \right\}. \quad (5)$$

- (2) These m local estimates are then averaged together—that is, we compute

$$\bar{\theta}_1 = \frac{1}{m} \sum_{i=1}^m \theta_{1,i}. \quad (6)$$

The bootstrap average mixture (BAVGM) algorithm is based on an additional level of sampling on top of the first, involving a fixed subsampling rate $r \in [0, 1]$. In particular, it consists of the following three steps:

Bootstrap average mixture algorithm

- (1) Each processor i draws a subset $S_{2,i}$ of size $\lceil rn \rceil$ by sampling uniformly at random without replacement from its local data set $S_{1,i}$.
- (2) Each processor i computes both the local empirical minimizers $\theta_{1,i}$ from equation (5), and the empirical minimizer

$$\theta_{2,i} \in \operatorname{argmin}_{\theta \in \Theta} \left\{ \underbrace{\frac{1}{|S_{2,i}|} \sum_{x \in S_{2,i}} f(\theta; x)}_{F_{2,i}(\theta)} \right\}. \quad (7)$$

- (3) In addition to the previous average (6), the BAVGM algorithm computes the bootstrap average $\bar{\theta}_2 := \frac{1}{m} \sum_{i=1}^m \theta_{2,i}$, and then returns the weighted combination

$$\bar{\theta}_{\text{BAVGM}} := \frac{\bar{\theta}_1 - r\bar{\theta}_2}{1 - r}. \quad (8)$$

The intuition for the weighted estimator (8) is similar to that for standard bias correction procedures using the bootstrap [9, 10]. Roughly speaking, if $b_0 = \theta^* - \theta_1$ is the bias of the first estimator, then we may approximate b_0 by the bootstrap estimate of bias $b_1 = \theta_1 - \theta_2$. Then, since $\theta^* = \theta_1 + b_0$, we use the fact that $b_1 \approx b_0$ to argue that $\theta^* = \theta_1 + b_0 \approx \theta_1 + b_1$. The re-normalization is needed to enforce that the relative “weights” of $\bar{\theta}_1$ and $\bar{\theta}_2$ sum to 1.

The goal of this paper is to understand under what conditions—and in what sense—the estimators (6) and (8) approach the *oracle performance*, by which we mean the error of a centralized risk minimization procedure that is given access to all $N = nm$ samples.

Notation: Before continuing, we define the remainder of our notation. We use ℓ_2 to denote the usual Euclidean norm $\|\theta\|_2 = (\sum_{j=1}^d \theta_j^2)^{\frac{1}{2}}$. The ℓ_2 -operator norm of a matrix $A \in \mathbb{R}^{d_1 \times d_2}$ is its maximum singular value, defined by

$$\|A\|_2 := \sup_{v \in \mathbb{R}^{d_2}, \|v\|_2 \leq 1} \|Av\|_2.$$

A convex function F is λ -strongly convex on a set $U \subseteq \mathbb{R}^d$ if for arbitrary $u, v \in U$ we have

$$F(u) \geq F(v) + \langle \nabla F(v), u - v \rangle + \frac{\lambda}{2} \|u - v\|_2^2.$$

(If F is not differentiable, we may replace ∇F with any subgradient of F .) We let \otimes denote the Kronecker product, and for a pair of vectors u, v , we define the outer product $u \otimes v = uv^\top$. For a three-times differentiable function F , we denote the third derivative tensor by $\nabla^3 F$, so that for each $u \in \operatorname{dom} F$ the operator $\nabla^3 F(u) : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$ is linear and satisfies the relation

$$[\nabla^3 F(u)(v \otimes v)]_i = \sum_{j,k=1}^d \left(\frac{\partial^3}{\partial u_i \partial u_j \partial u_k} F(u) \right) v_j v_k.$$

We denote the indicator function of an event \mathcal{E} by $1_{(\mathcal{E})}$, which is 1 if \mathcal{E} is true and 0 otherwise.

3 Theoretical results

Having described the AVGM and BAVGM algorithms, we now turn to describing their statistical properties, presenting our two main theorems in this section. After each theorem, we give some commentary and a few corollaries to help to better explain the behavior of the algorithms, as well as comparing our convergence results to previous work. To guarantee good estimation properties of our algorithms, we require regularity conditions on the empirical risks F_1 and F_2 . It is simplest to state these in terms of the sample functions $f(\theta; x)$, and we note that, as with Assumption B, these are only required to hold locally around the optimal point θ^* , in particular within some Euclidean ball $U = \{\theta \in \mathbb{R}^d \mid \|\theta^* - \theta\|_2 \leq \rho\} \subseteq \Theta$, where $\rho \in (0, 1)$ is the ball radius.

Assumption C (Smoothness). *There are finite constants G, H such that the first and the second partial derivatives of f exist and satisfy the bounds*

$$\mathbb{E}[\|\nabla f(\theta; X)\|_2^8] \leq G^8 \quad \text{and} \quad \mathbb{E}[\|\nabla^2 f(\theta; X) - \nabla^2 F_0(\theta)\|_2^8] \leq H^8 \quad \text{for all } \theta \in U.$$

In addition, for any $x \in \mathcal{X}$, the Hessian matrix $\nabla^2 f(\theta; x)$ is $L(x)$ -Lipschitz continuous, meaning that

$$\|\nabla^2 f(\theta'; x) - \nabla^2 f(\theta; x)\|_2 \leq L(x) \|\theta' - \theta\|_2 \quad \text{for all } \theta, \theta' \in U, \quad (9)$$

where, for some finite constant $L < \infty$, we have $\mathbb{E}[L(X)^8] \leq L^8$ and $\mathbb{E}[(L(X) - \mathbb{E}[L(X)])^8] \leq L^8$.

To be clear, it is actually necessary to impose some type of smoothness condition on the Hessian matrix, as in the Lipschitz condition (9), in order for averaging methods to work, as we now demonstrate. In fact, this example underscores the difficulty of proving that the AVGM algorithm achieves better mean-squared error than single-machine strategies.

Example (Necessity of Hessian conditions). Let X be a Bernoulli variable with parameter $\frac{1}{2}$, and consider the loss function

$$f(\theta; x) = \begin{cases} \theta^2 - \theta & \text{if } x = 0 \\ \theta^2 1_{(\theta \leq 0)} + \theta & \text{if } x = 1, \end{cases} \quad (10)$$

where $1_{(\theta \leq 0)}$ is the indicator function for the event $\{\theta \leq 0\}$. The associated population risk is $F_0(\theta) = \frac{1}{2}(\theta^2 + \theta^2 1_{(\theta \leq 0)})$. Since $|F'_0(w) - F'_0(v)| \leq 2|w - v|$, the population risk is strongly convex and smooth, but it has a discontinuous second derivative. By inspection, the unique minimizer of the population risk is $\theta^* = 0$, and by an asymptotic expansion we have that $\mathbb{E}[\theta_{1,i}] = \Omega(n^{-\frac{1}{2}})$. (See Appendix D for details of this step.) Consequently, the bias of $\bar{\theta}_1$ is $\Omega(n^{-\frac{1}{2}})$, and the AVGM algorithm using $N = mn$ observations must suffer mean squared error $\mathbb{E}[(\bar{\theta}_1 - \theta^*)^2] = \Omega(n^{-1})$. Thus, we see that some type of Hessian smoothness is necessary for fast rates.

Although the previous example establishes necessity, it is in some sense pathological: both the smoothness condition given in Assumption C and the local strong convexity condition given in Assumption B are relatively innocuous for practical problems. For instance, both conditions will hold for standard forms of regression, such as linear and logistic, as long as the *population* data covariance matrix is not rank deficient and the data has suitable moments. Moreover, in the linear regression case, one has $L = 0$.

3.1 Bounds for simple averaging

Our setup complete, we now turn to giving our theoretical results. Our first theorem gives the convergence rate of the AVGM procedure. We recall that θ^* denotes the minimizer of the population objective function F_0 , and that for each $i \in \{1, \dots, m\}$, we use S_i to denote a dataset of n independent samples. Recall that for each i , we use $\theta_i \in \arg \min_{\theta \in \Theta} \{\frac{1}{n} \sum_{x \in S_i} f(\theta; x)\}$ to denote a minimizer of the empirical risk for the dataset S_i , and we define the averaged vector $\bar{\theta} = \frac{1}{m} \sum_{i=1}^m \theta_i$. The following result bounds the mean-squared error between this averaged estimate and the minimizer θ^* of the population risk.

Theorem 1. *Under Assumptions A through C, there exists a finite numerical constant C such that*

$$\begin{aligned} \mathbb{E} \left[\|\bar{\theta} - \theta^*\|_2^2 \right] &\leq \frac{2}{nm} \mathbb{E} \left[\|\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; X)\|_2^2 \right] \\ &\quad + \frac{C}{\lambda^2 n^2} (H^2 \log d + L^2 G^2) \mathbb{E} \left[\|\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; X)\|_2^2 \right] \\ &\quad + \mathcal{O}(m^{-1} n^{-2}) + \mathcal{O}(n^{-3}). \end{aligned} \quad (11)$$

A simple (but somewhat weaker) corollary of Theorem 1 makes it easier to parse. In particular, note that

$$\|\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; x)\|_2 \stackrel{(i)}{\leq} \|\nabla^2 F_0(\theta^*)^{-1}\|_2 \|\nabla f(\theta^*; x)\|_2 \stackrel{(ii)}{\leq} \frac{1}{\lambda} \|\nabla f(\theta^*; x)\|_2, \quad (12)$$

where step (i) follows from the inequality $\|Ax\|_2 \leq \|A\| \|x\|_2$, valid for any matrix A and vector x ; and step (ii) follows from Assumption B. In addition, Assumption C implies $\mathbb{E}[\|\nabla f(\theta^*; X)\|_2^2] \leq G^2$, and so, putting together the pieces, we have established the following corollary.

Corollary 1. *Under the same conditions as Theorem 1,*

$$\mathbb{E} \left[\|\bar{\theta} - \theta^*\|_2^2 \right] \leq \frac{2G^2}{\lambda^2 nm} + \frac{CG^2}{\lambda^4 n^2} (H^2 \log d + L^2 G^2) + \mathcal{O}(m^{-1} n^{-2}) + \mathcal{O}(n^{-3}). \quad (13)$$

This upper bound shows that the leading term decays proportionally to $(nm)^{-1}$, with the pre-factor depending inversely on the strong convexity constant λ and growing proportionally with the bound G on the loss gradient. Although easily interpretable, the upper bound (13) can be loose, since it is based on the relatively weak bound (12)(i).

Note that the leading term in our original upper bound (11) involves the product of $\nabla f(\theta^*; X)$ with the inverse Hessian. In many statistical settings, including the problem of linear regression, the effect of this matrix-vector multiplication is to perform some type of standardization. When the loss $f(\cdot; x) : \Theta \rightarrow \mathbb{R}$ is actually the negative log-likelihood $\ell(x | \theta)$ for a parametric family of models $\{P_\theta\}$, we can make this intuition precise. In particular, under suitable regularity conditions (e.g., [15, Chapter 6]), we can define the Fisher information matrix

$$I(\theta^*) := \mathbb{E} \left[\nabla \ell(X | \theta^*) \nabla \ell(X | \theta^*)^\top \right] = \mathbb{E}[\nabla^2 \ell(X | \theta^*)].$$

Recalling that $N = mn$ is the total number of samples available, then under our assumptions, standard minimax theory guarantees for *any estimator* $\hat{\theta}_N$ based on N samples, essentially

$$\liminf_{N \rightarrow \infty} N \mathbb{E} \left[\|\hat{\theta}_N - \theta^*\|_2^2 \right] \geq \text{tr}(I(\theta^*)^{-1}).$$

For instance, see Theorem 8.11 in van der Vaart [29] for a result of this type. (A formal statement requires taking a type of supremum over the true parameter θ^* .) In connection with Theorem 1, we obtain:

Corollary 2. *In addition to the conditions of Theorem 1, suppose that the loss functions $f(\cdot; x)$ are the negative log-likelihood $\ell(x \mid \theta)$ for a parametric family $\{P_\theta, \theta \in \Theta\}$. Then there is a finite numerical constant C such that*

$$\mathbb{E} \left[\|\bar{\theta}_1 - \theta^*\|_2^2 \right] \leq \frac{2}{N} \text{tr}(I(\theta^*)^{-1}) + \frac{Cm^2 \text{tr}(I(\theta^*)^{-1})}{\lambda^2 N^2} (H^2 \log d + L^2 G^2) + \mathcal{O}(mN^{-2}).$$

Proof Rewriting the log-likelihood in the notation of Theorem 1, we have $\nabla \ell(x \mid \theta^*) = \nabla f(\theta^*; x)$ and all we need to note is that

$$\begin{aligned} I(\theta^*)^{-1} &= \mathbb{E} \left[I(\theta^*)^{-1} \nabla \ell(X \mid \theta^*) \nabla \ell(X \mid \theta^*)^\top I(\theta^*)^{-1} \right] \\ &= \mathbb{E} \left[(\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; X)) (\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; X))^\top \right]. \end{aligned}$$

Now apply the linearity of the trace and use the fact that $\text{tr}(uu^\top) = \|u\|_2^2$. □

Except for the factor of two in the bound, Corollary 2 shows that Theorem 1 essentially achieves the best possible result. The important aspect of our bound, however, is that we obtain this convergence rate without calculating an estimate on all $N = mn$ samples: instead, we calculate m independent estimators, and then average them to attain the convergence guarantee. We remark that an inspection of our proof shows that, at the expense of worse constants on higher order terms, we can reduce the factor of $2/mn$ on the leading term in Theorem 1 to $(1+c)/mn$ for any constant $c > 0$; as made clear by Corollary 2, this is essentially unimprovable.

3.2 Bounds for bootstrap mixture averaging

For small m , Theorem 1 and Corollary 1 show that the convergence rate of the AVGM algorithm is mainly determined by the first term in the bound (11), which is at most $\frac{G^2}{\lambda^2 mn}$. In contrast, when the number of processors m grows, the second term in the bound (11), in spite of being $\mathcal{O}(n^{-2})$, may have non-negligible effect. This issue will be exacerbated when the local strong convexity parameter λ of the risk F_0 is close to zero or the Lipschitz continuity constant H of $\nabla f(\theta; x)$ is large. This concern motivated our development of the bootstrap average mixture (BAVGM) algorithm, and we now turn to its analysis.

Due to the additional randomness introduced by the bootstrap algorithm BAVGM, its analysis requires an additional smoothness condition. In particular, recalling the Euclidean ρ -neighborhood U of the optimum θ^* , we require that the loss function f is (locally) smooth through its third derivatives.

Assumption D (Strong smoothness). *The third derivatives of f are Lipschitz continuous, meaning that for each $x \in \mathcal{X}$, there is a constant $M(x)$ such that*

$$\|(\nabla^3 f(\theta; x) - \nabla^3 f(\theta'; x)) (u \otimes u)\|_2 \leq M(x) \|\theta - \theta'\|_2 \|u\|_2^2 \quad \text{for all } \theta, \theta' \in U, \text{ and } u \in \mathbb{R}^d,$$

where $\mathbb{E}[M^4(X)] \leq M^4$ for some constant $M < \infty$.

It is easy to verify that Assumption D holds for least-squares regression (with $M = 0$); it also holds for logistic regression problems whose covariates have finite fourth moments.

We now state our second main theorem, which shows that the use of bootstrap samples to reduce the bias of the AVGM algorithm yields improved performance.

Theorem 2. *Under Assumptions A through D, the output $\bar{\theta}_{\text{BAVGM}} = (\bar{\theta}_1 - r\bar{\theta}_2)/(1 - r)$ of the bootstrap BAVGM algorithm has mean-squared error bounded as*

$$\begin{aligned} \mathbb{E} \left[\|\bar{\theta}_{\text{BAVGM}} - \theta^*\|_2^2 \right] &\leq \frac{2 + 3r}{(1 - r)^2} \cdot \frac{1}{nm} \mathbb{E} \left[\|\nabla^2 F_0(\theta^*)^{-1} \nabla f(\theta^*; X)\|_2^2 \right] \\ &\quad + \mathcal{O} \left(\frac{1}{(1 - r)^2} m^{-1} n^{-2} + \frac{1}{r(1 - r)^2} n^{-3} \right). \end{aligned}$$

Comparing the conclusions of Theorem 2 to those of Theorem 1, we see that the $\mathcal{O}(n^{-2})$ term in the bound (11) has been eliminated. The reason for this elimination is that subsampling at a rate r reduces the bias of the BAVGM algorithm to $\mathcal{O}(n^{-3})$; the bias of the AVGM algorithm induces terms of order n^{-2} in Theorem 1. Theorem 2 suggests that the performance of the BAVGM algorithm is affected by the subsampling rate r ; a rough estimate for r might be to take $r = \sqrt{m}/n$ so long as $m < n^2$. Roughly, when m becomes large we increase r , since the bias of the independent solutions may increase and we enjoy averaging affects from the BAVGM algorithm. When m is small, the BAVGM algorithm appears to provide limited benefits. We leave as an intriguing open question whether computing multiple bootstrap samples at each machine can yield improved performance or reduce the variance of the BAVGM procedure, and whether using estimates based on resampling the data with replacement—rather than subsampling without replacement—can yield improved performance.

3.3 Time complexity

In practice, the exact empirical minimizers assumed in Theorems 1 and 2 may be unavailable. Instead, we need to use a finite number of iterations of some optimization algorithm in order to obtain reasonable approximations to the exact minimizers. In this section, we sketch an argument that shows that both the AVGM algorithm and the BAVGM algorithm can use such approximate empirical minimizers, and as long as the optimization error is sufficiently small, the resulting averaged estimate achieves the same order-optimal statistical error. Here we provide the arguments only for the AVGM algorithm; the arguments for the BAVGM algorithm are completely similar.

More precisely, suppose that each processor runs a finite number of iterations of some optimization algorithm, thereby obtaining the vector θ'_i as an approximate minimizer of the objective function $F_{1,i}$, and hence an approximation to θ_i . Let $\bar{\theta}' = \frac{1}{m} \sum_{i=1}^m \theta'_i$ denote the average of these approximate minimizers, which corresponds to the output of the approximate AVGM algorithm. With this notation, we have

$$\mathbb{E} \left[\|\bar{\theta}' - \theta^*\|_2^2 \right] \stackrel{(i)}{\leq} 2\mathbb{E}[\|\bar{\theta} - \theta^*\|_2^2] + 2\mathbb{E}[\|\bar{\theta}' - \bar{\theta}\|_2^2] \stackrel{(ii)}{\leq} 2\mathbb{E}[\|\bar{\theta} - \theta^*\|_2^2] + 2\mathbb{E}[\|\theta'_1 - \theta_1\|_2^2]. \quad (14)$$

where step (i) follows by triangle inequality (and the elementary bound $(a + b)^2 \leq 2a^2 + 2b^2$), and step (ii) follows by Jensen's inequality. Consequently, suppose that each processor i runs enough iterations to obtain an approximate minimizer θ'_i such that

$$\mathbb{E}[\|\theta'_i - \theta_i\|_2^2] = \mathcal{O}((mn)^{-2}). \quad (15)$$

When this condition holds, the bound (14) shows that the average $\bar{\theta}'$ of the approximate minimizers shares the same convergence rates provided by Theorem 1.

We now consider how long it takes to compute an approximate minimizer θ'_i satisfying condition (15). In particular, assuming processing one sample requires one unit of time, we claim that this computation can be performed in time $\mathcal{O}(n \log(mn))$. In particular, the following two-stage strategy, involving a combination of stochastic gradient descent (see the following subsection for more details) and standard gradient descent, has this complexity:

- (1) First, as shown in the proof of Theorem 1, with high probability, the empirical risk F_1 is strongly convex in a ball $B_\rho(\theta_1)$ of constant radius $\rho > 0$ around θ_1 . Consequently, performing stochastic gradient descent on F_1 for $\mathcal{O}(\log^2(mn)/\rho^2)$ iterations yields an approximate minimizer that falls within $B_\rho(\theta_1)$ with high probability (e.g., see [20, Proposition 2.1]).
- (2) This initial estimate can be further improved by a few iterations of standard gradient descent. Under local strong convexity of the objective function, gradient descent is known to converge at a geometric rate [3, Chapter 9], so that $\mathcal{O}(\log(1/\epsilon))$ iterations will reduce the error to order ϵ . In our case, we have $\epsilon = (mn)^{-2}$, and since each iteration of standard gradient descent requires $\mathcal{O}(n)$ units of time, a total of $\mathcal{O}(n \log(mn))$ time units are sufficient to obtain a final estimate θ'_1 satisfying condition (15).

Overall, we conclude that the speed-up of the AVGM (and similar algorithms) over the naive approach of processing all $N = mn$ samples on one processor is at least of order $m/\log(N)$.

3.4 Stochastic gradient descent with averaging

The previous strategy involved a combination of stochastic gradient descent and standard gradient descent. In many settings, it may be appealing to use only a stochastic gradient algorithm, due to their ease of their implementation and limited computational requirements. In this section, we describe an extension of Theorem 1 to the case in which each machine computes an approximate minimizer using only stochastic gradient descent.

Stochastic gradient algorithms have a lengthy history in both statistics and optimization [27, 23, 20, 24]. Let us begin by briefly reviewing the basic form of stochastic gradient descent (SGD). Stochastic gradient descent algorithms iteratively update a parameter vector θ^t over time based on randomly sampled gradient information. Specifically, at iteration t , a sample X_t is drawn at random from the distribution P (or, in the case of a finite set of data $\{X_1, \dots, X_n\}$, a sample X_t is chosen from the data set). The method then performs the following two steps:

$$\theta^{t+\frac{1}{2}} = \theta^t - \eta_t \nabla f(\theta^t; X_t) \quad \text{and} \quad \theta^{t+1} = \operatorname{argmin}_{\theta \in \Theta} \left\{ \left\| \theta - \theta^{t+\frac{1}{2}} \right\|_2^2 \right\}. \quad (16)$$

Here $\eta_t > 0$ is a stepsize, and the first update in (16) is a gradient descent step with respect to the random gradient $\nabla f(\theta^t; X_t)$. The method then projects the intermediate point $\theta^{t+\frac{1}{2}}$ back onto the constraint set Θ (if there is a constraint set). The convergence of SGD methods of the form (16) has been well-studied, and we refer the reader to the papers [23, 20, 24] for deeper investigations.

To prove convergence of our stochastic gradient-based averaging algorithms, we require the following smoothness and strong convexity condition, which is an alternative to the Assumptions B and C used previously.

Assumption E (Smoothness and Strong Convexity II). *There exists a function $L : \mathcal{X} \rightarrow \mathbb{R}_+$ such that*

$$\|\nabla^2 f(\theta; x) - \nabla^2 f(\theta^*; x)\|_2 \leq L(x) \|\theta - \theta^*\|_2,$$

and $\mathbb{E}[L^2(X)] \leq L^2 < \infty$. There are finite constants G and H such that

$$\mathbb{E}[\|\nabla f(\theta; X)\|_2^4] \leq G^4, \quad \text{and} \quad \mathbb{E}[\|\nabla^2 f(\theta^*; X)\|_2^4] \leq H^4 \quad \text{for each fixed } \theta \in \Theta.$$

In addition, the population function F_0 is λ -strongly convex over the space Θ , meaning that

$$\nabla^2 F_0(\theta) \succeq \lambda I_{d \times d} \quad \text{for all } \theta \in \Theta. \quad (17)$$

Assumption E does not require as many moments as does Assumption C, but it does require each moment bound to hold globally, that is, over the entire sample space Θ , rather than only in a neighborhood of the optimal point θ^* . Similarly, the necessary curvature—in the form of the lower bound on the Hessian matrix $\nabla^2 F_0$ —is also required to hold globally, rather than only locally. Nonetheless, Assumption E holds for many common problems; for instance, it holds for any linear regression problem in which the covariates have finite fourth moments and the domain Θ is compact.

Moving to a description of the averaged stochastic gradient algorithm (SAVGM), it is based on the following two steps:

- (1) Given some constant $c > 1$, each machine performs n iterations of stochastic gradient descent (16) on its local dataset of n samples using the stepsize $\eta_t = \frac{c}{\lambda t}$, then outputs the resulting local parameter θ'_i .
- (2) The algorithm computes the average $\bar{\theta}^n = \frac{1}{m} \sum_{i=1}^m \theta'_i$.

The following result characterizes the mean-squared error of this procedure in terms of the constants

$$\alpha := \max \left\{ 4, \frac{c}{2 - 1/c} \right\} \quad \text{and} \quad \beta := \max \left\{ \left\lceil \frac{cH}{\lambda} \right\rceil, \frac{c\alpha^{3/4}G^{3/2}}{(c-1)\lambda^{5/2}} \left(\frac{\alpha^{1/4}LG^{1/2}}{\lambda^{1/2}} + \frac{4G + HR}{\rho^{3/2}} \right) \right\}.$$

Theorem 3. *Under Assumptions A and E, the output $\bar{\theta}^n$ of the SAVGM algorithm has mean-squared error upper bounded as*

$$\mathbb{E} \left[\|\bar{\theta}^n - \theta^*\|_2^2 \right] \leq \frac{\alpha G^2}{\lambda^2 mn} + \frac{\beta^2}{n^{3/2}}. \quad (18)$$

Theorem 3 shows that the averaged stochastic gradient descent procedure attains the optimal convergence rate as a function of the total number of observations $N = mn$. The constant and problem-dependent factors are somewhat worse than those in the earlier results we presented in Theorems 1 and 2, but the practical implementability of such a procedure may in some circumstances outweigh those differences. We also note that the second term of order $\mathcal{O}(n^{-3/2})$ may be reduced to $\mathcal{O}(n^{(2-2k)/k})$ for any $k \geq 4$ by assuming the existence of k th moments in Assumption E; we show this in passing after our proof of the theorem in Appendix C. It is not clear whether a bootstrap correction is possible for the stochastic-gradient based estimator; such a correction could be significant, because the term $\beta^2/n^{3/2}$ arising from the bias in the stochastic gradient estimator may be non-trivial. We leave this question to future work.

4 Performance on synthetic data

In this section, we report the results of simulation studies comparing the AVGM, BAVGM, and SAVGM methods, as well as a trivial method using only a fraction of the data available on a single machine. For each of our simulated experiments, we use a fixed total number of samples $N = 100,000$, but we vary the number of parallel splits m of the data (and consequently, the local dataset sizes $n = N/m$) and the dimensionality d of the problem solved.

For our experiments, we simulate data from one of two regression models:

$$y = \langle u, x \rangle + \varepsilon \quad \text{or} \quad (19a)$$

$$y = \langle u, x \rangle + h(x)|\varepsilon|, \quad (19b)$$

where $\varepsilon \sim N(0, 1)$, and h is a function to be specified. Specifically, the data generation procedure is as follows. For each individual simulation, we choose fixed vector $u \in \mathbb{R}^d$ with entries u_i distributed uniformly in $[0, 1]$. The data samples consist of pairs (x, y) , where $x \in \mathbb{R}^d$ and $y \in \mathbb{R}$ is the target value. To sample each x vector, we choose five distinct indices in $\{1, \dots, d\}$ uniformly at random, and the entries of x at those indices are distributed as $N(0, 1)$. In the sampling scheme (19a), the target y is sampled according to a standard normal regression model, so $\varepsilon \sim N(0, 1)$ is independent normal noise. We use the heteroskedastic sampling scheme (19b) to better understand the effects of bootstrap sampling in the BAVGM algorithm, as the standard least-squares estimator is unbiased for the model (19a). In the heteroskedastic model (19b), we still sample ε independently from $N(0, 1)$, but we set $h(x) = \sum_{i=1}^d (x_i/2)^3$, which is mean zero but forces the noise in y to be dependent on the data x .

Regardless of the sampling strategy, in our simulation experiments we use the least-squares loss

$$f(\theta; (x, y)) := \frac{1}{2}(\langle \theta, x \rangle - y)^2$$

to estimate the vector u . The goal in each experiment is then to estimate the vector θ^* that minimizes $F_0(\theta) := \mathbb{E}[f(\theta; (X, Y))]$. For each simulation, we generate N samples according to either the model (19a) or (19b). For each $m \in \{2, 4, 8, 16, 32, 64, 128\}$, we estimate $\theta^* = \operatorname{argmin}_{\theta} F_0(\theta)$ using a parallel method with data split into m independent sets of size $n = N/m$, specifically

- (i) The AVGM method
- (ii) The BAVGM method with several settings of the subsampling ratio r (we give specific values in the sequel)
- (iii) The SAVGM stochastic method using stepsize $\eta_t = d/(10(d+t))$, which gave empirically good performance.

In addition to (i)–(iii), we also estimate θ^* with

- (iv) The empirical minimizer of a single split of the data of size $n = N/m$
- (v) The empirical minimizer on the full dataset (the oracle solution).




Figure 1: The error $\|\hat{\theta} - \theta^*\|_2^2$ versus number of machines, with standard errors across twenty simulations, for solving least squares with data generated according to the normal model (19a). The oracle least-squares estimate using all N samples is given by the line “All,” while the line “Single” gives the performance of the naive estimator using only $n = N/m$ samples.




Figure 2: Comparison of AVGM and SAVGM methods as in Figure 1 plotted on logarithmic scale. The plot shows $\|\hat{\theta} - \theta^*\|_2^2 - \|\theta_N - \theta^*\|_2^2$, where θ_N is the oracle least-squares estimator using all N data samples.

4.1 Averaging methods

For our first set of experiments, we study the performance of the averaging methods (AVGM and SAVGM), showing their scaling as the number of splits of data—the number of machines m —grows for fixed N and dimensions $d = 20$ and $d = 200$. We use the standard regression model (19a) to generate the data, and throughout we let $\hat{\theta}$ denote the estimate returned by the method under




Figure 3: The error $\|\hat{\theta} - \theta^*\|_2^2$ versus number of machines, with standard errors across twenty simulations, for solving least squares with data generated according to the non-normal model (19b). The oracle least-squares estimate using all N samples is given by the line “All,” while the line “Single” gives the performance of the naive estimator using only $n = N/m$ samples.

consideration (so in the AVGM case, for example, this is the vector $\hat{\theta} := \bar{\theta}_1$). For the model (19a), the population optimal vector θ^* is simply u .

In Figure 1, we plot the error $\|\hat{\theta} - \theta^*\|_2^2$ of the inferred parameter vector $\hat{\theta}$ for the true parameters θ^* versus the number of splits m , or equivalently, the number of separate machines available for use. We also plot standard errors (across twenty experiments) for each curve. As a baseline in each plot, we plot as a red line the squared error $\|\hat{\theta}_N - \theta^*\|_2^2$ of the centralized “gold standard,” obtained by applying a batch method to all N samples.

From the plots in Figure 1, we can make a few observations. The AVGM algorithm enjoys excellent performance, as predicted by our theoretical results, especially compared to the naive solution using only a fraction $1/m$ of the data. In particular, if $\hat{\theta}$ is obtained by the batch method, then AVGM is almost as good as the full-batch baseline even for m as large as 128, though there is some evident degradation in solution quality. The SAVGM (stochastic-gradient with averaging) solution also yields much higher accuracy than the naive solution, but its performance degrades somewhat more quickly than the AVGM method’s as m grows. In higher dimensions, both the AVGM and SAVGM procedures have somewhat worse performance; again, this is not unexpected.

We present a comparison between the AVGM method and the SAVGM method with somewhat more distinguishing power in Figure 2. For these plots, we compute the gap between the AVGM mean-squared-error and the unparallel baseline MSE, which is the accuracy lost due to parallelization or distributing the inference procedure across multiple machines. Figure 2 shows that the mean-squared error grows polynomially with the number of machines m , which is consistent with our theoretical results. Corollary 2 shows that we expect the AVGM method to suffer (lower-order) penalties proportional to m^2 as m grows, while Theorem 3 suggests the somewhat faster growth we see for the SAVGM method in Figure 2. Thus we see that the improved run-time performance of the SAVGM method—it requires only a single pass through the data on each machine, touching each datum only once—comes at the expense of some loss of accuracy, as measured by mean-squared error.




Figure 4: The error $\|\hat{\theta} - \theta^*\|_2^2$ plotted against the number of machines m for the AVGM and BAVGM methods, with standard errors across twenty simulations, using the normal regression model (19a). The oracle estimator is denoted by the line “All.”

In Figure 3, we repeat the simulations generating Figure 1, but we use the non-normal regression model (19b) to generate the data. In this case, the non-normality of the model means that the vector u generating the data is no longer equal to θ^* , so we estimate θ^* by solving a least-squares regression problem with $10N$ the number of data points, and use this to measure the mean-squared error $\|\hat{\theta} - \theta^*\|_2^2$. Broadly, we see that the standard averaging (AVGM) and stochastic (SAVGM) methods provide a strong benefit over naive solutions, achieving performance close to the gold-standard oracle estimator using all the data.

4.2 Bootstrap correction

We now turn to developing an understanding of the BAVGM algorithm in comparison to the standard average mixture algorithm, developing intuition for the benefits and drawbacks of the method. Before describing the results, we remark that for the standard regression model (19a), the least-squares solution is unbiased for θ^* , so we expect bootstrap averaging to yield little (if any) improvements. The BAVGM method is essentially aimed at correcting the bias of the estimator $\bar{\theta}_1$, and de-biasing an unbiased estimator only increases its variance. However, for the non-normal model (19b) we do expect to see some performance gains. In our experiments, we use multiple sub-sampling rates to study their effects, choosing $r \in \{0.005, 0.01, 0.02, 0.04\}$. We remind the reader that the output of the BAVGM algorithm is the vector $\hat{\theta} := (\bar{\theta}_1 - r\bar{\theta}_2)/(1 - r)$.

In Figure 4, we plot the results of simulations comparing AVGM and BAVGM with data generated from the normal regression model (19a). Both algorithms have low error rates, but the AVGM method is slightly better than the BAVGM method for both values of the dimension d and all and sub-sampling rates r . As expected, in this case the BAVGM method does not offer improvement over AVGM. (In Figure 4(a), we note that the standard error is in fact very small, but the mean-squared error is only of order 10^{-3} .)

To that end, in Figure 5, we plot analogous mean-square error curves for the least-squares re-

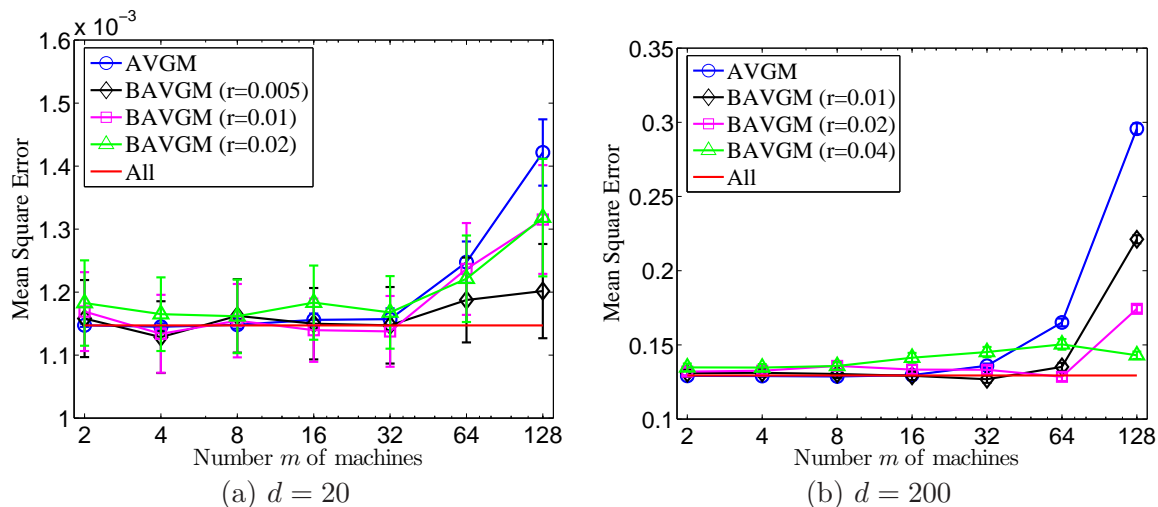


Figure 5: The error $\|\hat{\theta} - \theta^*\|_2^2$ plotted against the number of machines m for the AVGM and BAVGM methods, with standard errors across twenty simulations, using the non-normal regression model (19b). The oracle estimator is denoted by the line “All.”

gression problem when the vector y is sampled according to the non-normal regression model (19b). In this case, the least-squares estimator is biased for θ^* (which, as before, we estimate by solving a larger regression problem using $10N$ data samples). Figure 5 shows that both the AVGM and BAVGM method still enjoy good performance; in some cases, the BAVGM method even beats the oracle least-squares estimator for θ^* that uses all N samples. Both are clearly better than the naive estimator based on N/m samples (recall Figure 3). Since the AVGM estimate is biased in this case, its error curve increases roughly quadratically with m , which agrees with our theoretical predictions in Theorem 1. In contrast, we see that the BAVGM algorithm enjoys somewhat more stable performance, with increasing benefit as the number of machines m increases. For example, in case of $d = 200$, if we choose $r = 0.01$ for $m \leq 32$, choose $r = 0.02$ for $m = 64$ and $r = 0.04$ for $m = 128$, then BAVGM has performance comparable with the oracle method that uses all N samples. Moreover, we see that all the values of r —at least for the reasonably small values we use in the experiment—provide performance improvements over a non-bootstrapped distributed estimator.

5 Experiments with advertising data

Predicting whether a user of a search engine will click on an advertisement presented to him or her is of central importance to the business of several internet companies, and in this section, we present experiments studying the performance of the AVGM and BAVGM methods for this task. We use a large dataset from the Tencent search engine, `soso.com` [28], which contains 641,707 distinct advertisement items with $N = 235,582,879$ data samples.

Each sample consists of a so-called *impression*, which in the terminology of the information retrieval literature (e.g., see the book by Manning et al. [17]), is a list containing a user-issued search, the advertisement presented to the user in response to the search, and a label $y \in \{+1, -1\}$ indicating whether the user clicked on the advertisement. The ads in our dataset were presented

Table 1: Features used in online advertisement prediction problem.

Feature Name	Dimension	Description
Query	20000	Word tokens appearing in the query.
Gender	3	Gender of the user
Keyword	20000	Word tokens appearing in the purchase keywords.
Title	20000	Word tokens appearing in the ad title.
Advertiser	39191	Advertiser’s ID
AdID	641707	Advertisement’s ID.
Age	6	Age of the user
UserFreq	25	Number of appearance of the same user.
Position	3	Position of advertisement on search page.
Depth	3	Number of ads in the session.
QueryFreq	25	Number of occurrences of the same query.
AdFreq	25	Number of occurrences of the same ad.
QueryLength	20	Number of words in the query.
TitleLength	30	Number of words in the ad title.
DespLength	50	Number of words in the ad description.
QueryCtr	150	Average click-through-rate for query.
UserCtr	150	Average click-through-rate for user.
AdvrCtr	150	Average click-through-rate for advertiser.
WordCtr	150	Average click-through-rate for keyword advertised.
UserAdFreq	20	Number of times this user sees an ad.
UserQueryFreq	20	Number of times this user performs a search.

to 23,669,283 distinct users.

Transforming an impression into a useable set of regressors x is non-trivial, but the Tencent dataset provides a standard encoding. We list the features present in the data in Table 1, along with some description of their meaning. Each text-based feature—that is, those made up of words, which are Query, Keyword, and Title—is given a “bag-of-words” encoding [17]. This encoding assigns each of 20,000 possible words an index, and if the word appears in the query (or Keyword or Title feature), the corresponding index in the vector x is set to 1. Words that do not appear are encoded with a 0. Real-valued features, corresponding to the bottom fifteen features in Table 1 beginning with “Age”, are binned into a fixed number of intervals $[-\infty, a_1], (a_1, a_2], \dots, (a_k, \infty]$, each of which is assigned an index in x . (Note that the intervals and number thereof vary per feature, and the dimension of the features listed in Table 1 corresponds to the number of intervals). When a feature falls into a particular bin, the corresponding entry of x is assigned a 1, and otherwise the entries of x corresponding to the feature are 0. Each feature has one additional value for “unknown.” The remaining categorical features—gender, advertiser, and advertisement ID (AdID)—are also given $\{0, 1\}$ encodings, where only one index of x corresponding to the feature may be non-zero (which indicates the particular gender, advertiser, or AdID). This combination of encodings yields a binary-valued covariate vector $x \in \{0, 1\}^d$ with $d = 741,725$ dimensions.

Our goal is to predict the probability of a user clicking a given advertisement as a function of the covariates in Table 1. In order to do so, we use a logistic regression model to estimate the

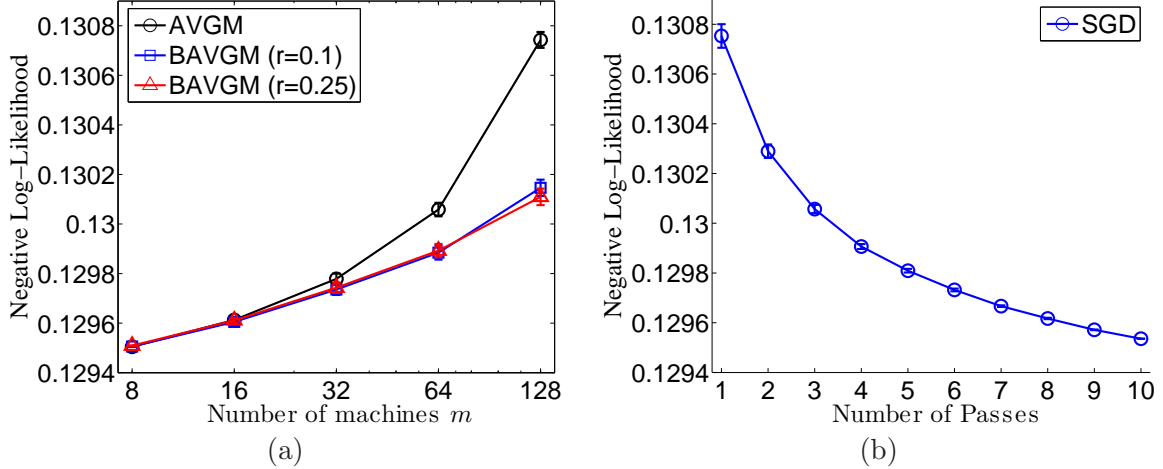


Figure 6: The negative log-likelihood of the output of the AVGM, BAVGM, and a stochastic gradient method (the update (16)) on the held-out dataset for the click-through prediction task. (a) Performance of the AVGM and BAVGM methods versus the number of splits m of the data. (b) Performance of the SGD baseline as a function of number of passes through the entire dataset.

probability of a click response

$$P(y = 1 \mid x; \theta) := \frac{1}{1 + \exp(-\langle \theta, x \rangle)},$$

where $\theta \in \mathbb{R}^d$ is the unknown regression vector. We use the negative logarithm of P as the loss, incorporating a ridge regularization penalty. This combination yields the overall optimization problem

$$f(\theta; (x, y)) = \log(1 + \exp(-y \langle \theta, x \rangle)) + \frac{\lambda}{2} \|\theta\|_2^2. \quad (20)$$

In all our experiments, we use regularization parameter $\lambda = 10^{-6}$, a choice obtained by cross validation.

For this problem, we cannot evaluate the mean-squared error $\|\hat{\theta} - \theta^*\|_2^2$, as we do not know the true optimal parameter θ^* . Consequently, we evaluate the performance of an estimate $\hat{\theta}$ using log-loss on a held-out dataset. Specifically, we perform a five-fold validation experiment, where we shuffle the data and partition it into five equal-sized subsets. For each of our five experiments, we hold out one partition to use as the test set, using the remaining data as the training set for inference. When studying the AVGM or BAVGM method, we compute the local estimate θ_i via a trust-region Newton-based method [21].

The dataset is too large to fit in the memory of most computers: in total, four splits of the data require 55 gigabytes. Consequently, it is difficult to provide an oracle training comparison using the full N samples. Instead, for each experiment, we perform 10 passes of stochastic gradient descent through the dataset to get a rough baseline of the performance attained by the empirical minimizer for the entire training dataset. Figure 6(b) shows the hold-out set log-loss after each of the sequential passes through the training data finishes.

In Figure 6(a), we show the average hold-out set log-loss (with standard errors) of the estimator $\bar{\theta}_1$ provided by the AVGM method versus number of splits of the data m , and we also plot the log-loss of the BAVGM method using subsampling ratios of $r \in \{.1, .25\}$. The plot shows that for small




Figure 7: The log-loss on held-out data for the BAVGM method applied with $m = 128$ parallel splits of the data, plotted versus the sub-sampling rate r .

m , both AVGM and BAVGM enjoy good performance, comparable to or better than (our proxy for) the oracle solution using all N samples. As the number of machines m grows, however, the de-biasing provided by the subsampled bootstrap method yields substantial improvements over the standard AVGM method. In addition, even with $m = 128$ splits of the dataset, the BAVGM method gives better hold-out set performance than performing two passes of stochastic gradient on the entire dataset of m samples; with $m = 64$, BAVGM enjoys performance as strong as looping through the data four times with stochastic gradient descent. This is striking, as doing even one pass through the data with stochastic gradient descent is known to give minimax optimal convergence rates [23, 2].

It is instructive and important to understand the sensitivity of the BAVGM method to the value of the resampling parameter r . We explore this question in Figure 7 using $m = 128$ splits, where we plot the log-loss of the BAVGM estimator on the held-out data set versus the subsampling ratio r . We choose $m = 128$ because more data splits provide more variable performance in r . For the *soso.com* ad prediction data set, the choice $r = .25$ achieves the best performance, but Figure 7 suggests that mis-specifying the ratio is not terribly detrimental. Indeed, while the performance of BAVGM degrades to that of the AVGM method, a wide range of settings of r give improved performance, and there does not appear to be a phase transition to poor performance.

6 Discussion

Large scale statistical inference problems are challenging, and the difficulty of solving them will only grow as data becomes more abundant: the amount of data we collect is growing much faster than the speed or storage capabilities of our computers. Our AVGM and BAVGM methods provide strategies for efficiently solving such large-scale risk minimization problems, enjoying performance comparable to an oracle method that is able to access the entire large dataset. We believe there are several interesting questions that remain open after this work. First, reproducing kernel Hilbert space methods, which suffer scaling cubic in the size of the data, may benefit from the distributed strategies in this paper, though their analysis is not immediate. More generally, an understanding of the interplay between statistical efficiency and communication could provide an avenue for further

research, and it may also be interesting to study the effects of subsampled or bootstrap-based estimators in other distributed environments.

Acknowledgments

We thank Joel Tropp for some informative discussions on and references for matrix concentration and moment inequalities.

A Proof of Theorem 1

Although Theorem 1 is in terms of bounds on 8^{th} order moments, we prove a somewhat more general result in terms of a set of (k_0, k_1, k_2) moment conditions given by

$$\begin{aligned}\mathbb{E}[\|\nabla f(\theta; X)\|_2^{k_0}] &\leq G^{k_0}, & \mathbb{E}[\|\nabla^2 f(\theta; X) - \nabla^2 F_0(\theta)\|_2^{k_1}] &\leq H^{k_1}, \\ \mathbb{E}[L(X)^{k_2}] &\leq L^{k_2} \quad \text{and} \quad \mathbb{E}[(L(X) - \mathbb{E}[L(X)])^{k_2}] &\leq L^{k_2}\end{aligned}$$

for $\theta \in U$. (Recall the definition of U prior Assumption C). Doing so allows sharper control if higher moments are known on certain quantities. The reader should recall throughout our arguments that we have assumed $\min\{k_0, k_1, k_2\} \geq 8$. Throughout the proof, we use F_1 and θ_1 to indicate the local empirical objective and empirical minimizer of machine 1 (which have the same distribution as those of the other processors), and we recall the notation $1_{\mathcal{E}}$ for the indicator function of the event \mathcal{E} .

Before beginning the proof of Theorem 1 proper, we begin with a simple inequality that relates the error term $\bar{\theta} - \theta^*$ to an average of the errors $\theta_i - \theta^*$, each of which we can bound in turn. Specifically, a bit of algebra gives us that

$$\begin{aligned}\mathbb{E}[\|\bar{\theta} - \theta^*\|_2^2] &= \mathbb{E}\left[\left\|\frac{1}{m} \sum_{i=1}^m \theta_i - \theta^*\right\|_2^2\right] \\ &= \frac{1}{m^2} \sum_{i=1}^m \mathbb{E}[\|\theta_i - \theta^*\|_2^2] + \frac{1}{m^2} \sum_{i \neq j} \mathbb{E}[\langle \theta_i - \theta^*, \theta_j - \theta^* \rangle] \\ &= \frac{1}{m} \mathbb{E}[\|\theta_1 - \theta^*\|_2^2] + \frac{m(m-1)}{m^2} \|\mathbb{E}[\theta_1 - \theta^*]\|_2^2 \\ &\leq \frac{1}{m} \mathbb{E}[\|\theta_1 - \theta^*\|_2^2] + \|\mathbb{E}[\theta_1 - \theta^*]\|_2^2.\end{aligned}\tag{21}$$

Here we used the definition of the averaged vector $\bar{\theta}$ and the fact that for $i \neq j$, the vectors θ_i and θ_j are independent (they are functions of independent samples). The upper bound (21) illuminates the path for the remainder of our proof: we bound each of $\mathbb{E}[\|\theta_i - \theta^*\|_2^2]$ and $\|\mathbb{E}[\theta_i - \theta^*]\|_2^2$. Intuitively, since our objective is (locally) strongly convex by Assumption B, the empirical minimizing vector θ_1 is a nearly unbiased estimator for θ^* , which allows us to prove the convergence rates in the theorem.

We begin by defining three events—which we (later) show hold with high probability—that guarantee the closeness of θ_1 and θ^* . Intuitively, these events imply that the function F_1 looks quite similar to the population risk F_0 around the point θ^* ; since F_0 is (locally) strongly convex,

the minimizer θ_1 of F_1 will be close to θ^* . Recall that Assumption C guarantees the existence of a ball $U_\rho = \{\theta \in \mathbb{R}^d : \|\theta - \theta^*\|_2 < \rho\}$ of radius $\rho \in (0, 1)$ such that

$$\|\nabla^2 f(\theta; x) - \nabla^2 f(\theta'; x)\|_2 \leq L(x) \|\theta - \theta'\|_2$$

for all $\theta, \theta' \in U_\rho$ and any x , where $\mathbb{E}[L(X)^{k_2}] \leq L^{k_2}$. In addition, Assumption B guarantees that $\nabla^2 F_0(\theta^*) \succeq \lambda I$. Now, choosing the potentially smaller radius $\delta_\rho = \min\{\rho, \rho\lambda/4L\}$, we can define the three “good” events

$$\begin{aligned} \mathcal{E}_0 &:= \left\{ \frac{1}{n} \sum_{i=1}^n L(X_i) \leq 2L \right\}, \\ \mathcal{E}_1 &:= \left\{ \|\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*)\|_2 \leq \frac{\rho\lambda}{2} \right\}, \quad \text{and} \\ \mathcal{E}_2 &:= \left\{ \|\nabla F_1(\theta^*)\|_2 \leq \frac{(1-\rho)\lambda\delta_\rho}{2} \right\}. \end{aligned} \tag{22}$$

We then have the following lemma, which states that as long as F_0 and F_1 are similar near their optima, then the distance between $\theta_1 \in \operatorname{argmin}_{\theta \in \Theta} F_1(\theta)$ and θ^* decreases with the norm of the gradient $\|\nabla F_1(\theta^*)\|_2$: more precisely, so long as $\nabla F_1(\theta^*)$ is small, then $\|\theta_1 - \theta^*\|_2$ decreases at least linearly in $\|\nabla F_1(\theta^*)\|_2$.

Lemma 1. *Under the events \mathcal{E}_0 , \mathcal{E}_1 , and \mathcal{E}_2 previously defined (22), we have*

$$\|\theta_1 - \theta^*\|_2 \leq \frac{2\|\nabla F_1(\theta^*)\|_2}{(1-\rho)\lambda}, \quad \text{and} \quad \nabla^2 F_1(\theta) \succeq (1-\rho)\lambda I_{d \times d}.$$

The proof of Lemma 1 relies on some standard optimization guarantees relating gradients to minimizers of functions (e.g. [3], Chapter 9), although some care is required since smoothness and strong convexity hold only locally in our problem. As the argument is somewhat technical, we defer it to Appendix E.

Our approach from here is to give bounds on $\mathbb{E}[\|\theta_1 - \theta^*\|_2^2]$ and $\|\mathbb{E}[\theta_1 - \theta^*]\|_2^2$ by careful Taylor expansions, which allows us to bound $\mathbb{E}[\|\bar{\theta}_1 - \theta^*\|_2^2]$ via our initial expansion (21). Similar Taylor expansions are of course familiar from asymptotic convergence in distribution results for M -estimators (e.g. [13], Chapter 9), but we require more care to control a larger number of moments. We begin by noting that whenever the events \mathcal{E}_0 , \mathcal{E}_1 , and \mathcal{E}_2 hold, then $\nabla F_1(\theta_1) = 0$, and moreover, by a Taylor series expansion of ∇F_1 between θ^* and θ_1 , we have

$$0 = \nabla F_1(\theta_1) = \nabla F_1(\theta^*) + \nabla^2 F_1(\theta')(\theta_1 - \theta^*)$$

where $\theta' = \kappa\theta^* + (1-\kappa)\theta_1$ for some $\kappa \in [0, 1]$. By adding and subtracting terms, we have

$$\begin{aligned} 0 &= \nabla F_1(\theta^*) + (\nabla^2 F_1(\theta') - \nabla^2 F_1(\theta^*))(\theta_1 - \theta^*) \\ &\quad + (\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*))(\theta_1 - \theta^*) + \nabla^2 F_0(\theta^*)(\theta_1 - \theta^*). \end{aligned} \tag{23}$$

Since $\nabla^2 F_0(\theta^*) \succeq \lambda I$, we can define the inverse Hessian matrix $\Sigma^{-1} := [\nabla^2 F_0(\theta^*)]^{-1}$, and setting $\Delta := \theta_1 - \theta^*$, we multiply both sides of the Taylor expansion (23) by Σ^{-1} to obtain the relation

$$\Delta = -\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(\nabla^2 F_1(\theta^*) - \nabla^2 F_1(\theta'))\Delta + \Sigma^{-1}(\nabla^2 F_0(\theta^*) - \nabla^2 F_1(\theta^*))\Delta. \tag{24}$$

Thus, if we define the matrices $P = \nabla^2 F_0(\theta^*) - \nabla^2 F_1(\theta^*)$ and $Q = \nabla^2 F_1(\theta^*) - \nabla^2 F_1(\theta')$, equality (24) can be re-written as

$$\theta_1 - \theta^* = -\Sigma^{-1} \nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*). \quad (25)$$

Note that equation (25) holds when the conditions of Lemma 1 hold, and otherwise we may simply assert only that $\|\theta_1 - \theta^*\|_2 \leq R$. Roughly, we expect the final two terms in the error expansion (25) to be of smaller order than the first term, since we hope that $\theta_1 - \theta^* \rightarrow 0$ and additionally that the Hessian differences decrease to zero at a sufficiently fast rate. We now formalize this intuition.

Inspecting the Taylor expansion (25), we see that there are several terms of a form similar to $(\nabla^2 F_0(\theta^*) - \nabla^2 F_1(\theta^*))(\theta_1 - \theta^*)$; using the smoothness Assumption C, we can convert these terms into higher order terms involving only $\theta_1 - \theta^*$. Thus, to effectively control the expansions (24) and (25), we must show that higher order terms of the form $\mathbb{E}[\|\theta_1 - \theta^*\|_2^k]$, for $k \geq 2$, decrease quickly enough in n .

Control of $\mathbb{E}[\|\theta_1 - \theta^*\|_2^k]$: Recalling the events (22), we define $\mathcal{E} := \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$ and then observe that

$$\begin{aligned} \mathbb{E}[\|\theta_1 - \theta^*\|_2^k] &= \mathbb{E}[1_{(\mathcal{E})} \|\theta_1 - \theta^*\|_2^k] + \mathbb{E}[1_{(\mathcal{E}^c)} \|\theta_1 - \theta^*\|_2^k] \\ &\leq \frac{2^k \mathbb{E}[1_{(\mathcal{E})} \|\nabla F_1(\theta^*)\|_2^k]}{(1 - \rho)^k \lambda^k} + \mathbb{P}(\mathcal{E}^c) R^k \\ &\leq \frac{2^k \mathbb{E}[\|\nabla F_1(\theta^*)\|_2^k]}{(1 - \rho)^k \lambda^k} + \mathbb{P}(\mathcal{E}^c) R^k, \end{aligned}$$

where we have used the bound $\|\theta - \theta^*\|_2 \leq R$ for all $\theta \in \Theta$, from Assumption A. Our goal is to prove that $\mathbb{E}[\|\nabla F_1(\theta^*)\|_2^k] = \mathcal{O}(n^{-k/2})$ and that $\mathbb{P}(\mathcal{E}^c) = \mathcal{O}(n^{-k/2})$.

We move forward with a few somewhat technical (but short) lemmas that lay the groundwork for proving these two facts.

Lemma 2. *Under Assumption C, there exist constants C and C' (dependent only on the moments k_0 and k_1 respectively) such that*

$$\mathbb{E}[\|\nabla F_1(\theta^*)\|_2^{k_0}] \leq C \frac{G^{k_0}}{n^{k_0/2}}, \quad \text{and} \quad (26a)$$

$$\mathbb{E}[\|\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*)\|^{k_1}] \leq C' \frac{\log^{k_1/2}(2d)H^{k_1}}{n^{k_1/2}}. \quad (26b)$$

See Appendix F for the proof of this claim. As an immediate consequence of Lemma 2, we see that the events \mathcal{E}_1 and \mathcal{E}_2 defined by (22) occur reasonably high probability. Indeed, recalling that $\mathcal{E} = \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$, Boole's law and the union bound imply

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &= \mathbb{P}(\mathcal{E}_0^c \cup \mathcal{E}_1^c \cup \mathcal{E}_2^c) \\ &\leq \mathbb{P}(\mathcal{E}_0^c) + \mathbb{P}(\mathcal{E}_1^c) + \mathbb{P}(\mathcal{E}_2^c) \\ &\leq \frac{\mathbb{E}[(\frac{1}{n} \sum_{i=1}^n L(X_i) - \mathbb{E}[L(X)])^{k_2}]}{L^{k_2}} + \frac{2^{k_1} \mathbb{E}[\|\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*)\|^{k_1}]}{\rho^{k_1} \lambda^{k_1}} + \frac{2^{k_0} \mathbb{E}[\|\nabla F_1(\theta^*)\|_2^{k_0}]}{(1 - \rho)^{k_0} \lambda^{k_0} \delta_\rho^{k_0}} \\ &\leq C_2 \frac{1}{n^{k_2/2}} + C_1 \frac{\log^{k_1/2}(2d)H^{k_1}}{n^{k_1/2}} + C_0 \frac{G^{k_0}}{n^{k_0/2}} \end{aligned} \quad (27)$$

for some universal constants C_0, C_1, C_2 . Consequently, we find that $\mathbb{P}(\mathcal{E}^c)R^k = \mathcal{O}(R^k(n^{-k_1/2} + n^{-k_2/2} + n^{-k_0/2}))$ for any $k \in \mathbb{N}$. In conclusion, we have proved the following lemma:

Lemma 3. *Let Assumptions B and C hold. For any $k \in \mathbb{N}$ with $k \leq \min\{k_0, k_1, k_2\}$, we have*

$$\mathbb{E}[\|\theta_1 - \theta^*\|_2^k] = \mathcal{O}\left(n^{-k/2} \cdot \frac{G^k}{(1-\rho)^k \lambda^k} + n^{-k_0/2} + n^{-k_1/2} + n^{-k_2/2}\right) = \mathcal{O}\left(n^{-k/2}\right),$$

where the order statements hold as $n \rightarrow +\infty$.

Recall the matrix $Q = \nabla^2 F_1(\theta^*) - \nabla^2 F_1(\theta')$ defined following equation (24). The following result controls the moments of its operator norm:

Lemma 4. *For $k \leq \min\{k_2, k_1, k_0\}/2$, we have $\mathbb{E}[\|Q\|_2^k] = \mathcal{O}(n^{-k/2})$.*

Proof We begin by using Jensen's inequality and Assumption C to see that

$$\|Q\|^k \leq \frac{1}{n} \sum_{i=1}^n \|\nabla^2 f(\theta'; X_i) - \nabla^2 f(\theta^*; X_i)\|^k \leq \frac{1}{n} \sum_{i=1}^n L(X_i)^k \|\theta' - \theta^*\|_2^k.$$

Now we apply the Cauchy-Schwarz inequality and Lemma 3, thereby obtaining

$$\mathbb{E}[\|Q\|_2^k] \leq \mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^n L(X_i)^k\right)^2\right]^{\frac{1}{2}} \mathbb{E}[\|\theta_1 - \theta^*\|_2^{2k}]^{\frac{1}{2}} = \mathcal{O}\left(L^k \frac{G^k}{(1-\rho)^k \lambda^k} n^{-k/2}\right),$$

where we have used Assumption C again. □

Lemma 3 allows us to control the first term from our initial bound (21) almost immediately. Indeed, using our last Taylor expansion (25) and the definition of the event $\mathcal{E} = \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$, we have

$$\begin{aligned} \mathbb{E}[\|\theta_1 - \theta^*\|_2^2] &= \mathbb{E}\left[1_{(\mathcal{E})} \|\Sigma^{-1} \nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*)\|_2^2\right] + \mathbb{E}[1_{(\mathcal{E}^c)} \|\theta_1 - \theta^*\|_2^2] \\ &\leq 2\mathbb{E}\left[\|\Sigma^{-1} \nabla F_1(\theta^*)\|_2^2\right] + 2\mathbb{E}\left[\|\Sigma^{-1}(P + Q)(\theta_1 - \theta^*)\|_2^2\right] + \mathbb{P}(\mathcal{E}^c)R^2, \end{aligned}$$

where we have applied the inequality $(a + b)^2 \leq 2a^2 + 2b^2$. Again using this same inequality, then applying Cauchy-Schwarz and Lemmas 2 and 4, we see that

$$\begin{aligned} \mathbb{E}\left[\|\Sigma^{-1}(P + Q)(\theta_1 - \theta^*)\|_2^2\right] &\leq 2\|\Sigma^{-1}\|_2^2 \left(\mathbb{E}[\|P\|_2^2 \|\theta_1 - \theta^*\|_2^2] + \mathbb{E}[\|Q\|_2^2 \|\theta_1 - \theta^*\|_2^2]\right) \\ &\leq 2\|\Sigma^{-1}\|_2^2 \left(\sqrt{\mathbb{E}[\|P\|_2^4] \mathbb{E}[\|\theta_1 - \theta^*\|_2^4]} + \sqrt{\mathbb{E}[\|Q\|_2^4] \mathbb{E}[\|\theta_1 - \theta^*\|_2^4]}\right) \\ &= \mathcal{O}(n^{-2}), \end{aligned}$$

where we have used the fact that $\min\{k_0, k_1, k_2\} \geq 8$ to apply Lemma 4. Combining these results, we obtain the upper bound

$$\mathbb{E}[\|\theta_1 - \theta^*\|_2^2] \leq 2\mathbb{E}\left[\|\Sigma^{-1} \nabla F_1(\theta^*)\|_2^2\right] + \mathcal{O}(n^{-2}), \quad (28)$$

which completes the first part of our proof of Theorem 1.

Control of $\|\mathbb{E}[\theta_1 - \theta^*]\|_2^2$: It remains to consider the $\|\mathbb{E}[\theta_1 - \theta^*]\|_2^2$ term from our initial error inequality (21). When the events (22) occur, we know that all derivatives exist, so we may recursively apply our expansion (25) of $\theta_1 - \theta^*$ to find that

$$\begin{aligned}\theta_1 - \theta^* &= -\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*) \\ &= \underbrace{-\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q) [-\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*)]}_{=:v}\end{aligned}\quad (29)$$

where we have introduced v as shorthand for the vector on the right hand side. Thus, with a bit of algebraic manipulation we obtain the relation

$$\theta_1 - \theta^* = 1_{(\mathcal{E})}v + 1_{(\mathcal{E}^c)}(\theta_1 - \theta^*) = v + 1_{(\mathcal{E}^c)}(\theta_1 - \theta^*) - 1_{(\mathcal{E}^c)}v = v + 1_{(\mathcal{E}^c)}(\theta_1 - \theta^* - v). \quad (30)$$

Now note that

$$\begin{aligned}\mathbb{E}[v] &= \mathbb{E}[-\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)[- \Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(w_1 - w^*)]] \\ &= \mathbb{E}[\Sigma^{-1}(P + Q)[\Sigma^{-1}(P + Q)(\theta_1 - \theta^*) - \nabla F_1(\theta^*)]].\end{aligned}$$

Thus, by re-substituting the appropriate quantities in (30) and applying the triangle inequality, we have

$$\begin{aligned}\|\mathbb{E}[\theta_1 - \theta^*]\|_2 &\leq \|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}((P + Q)(w_1 - w^*) - \nabla F_1(\theta^*))]\|_2 + \|\mathbb{E}[1_{(\mathcal{E}^c)}(\theta_1 - \theta^* - v)]\|_2 \\ &\leq \|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}((P + Q)(w_1 - w^*) - \nabla F_1(\theta^*))]\|_2 + \mathbb{E}[1_{(\mathcal{E}^c)}\|\theta_1 - \theta^*\|_2] \\ &\quad + \mathbb{E}[1_{(\mathcal{E}^c)}\|-\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)[- \Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*)]\|_2].\end{aligned}\quad (31)$$

Since $\|\theta_1 - \theta^*\|_2 \leq R$ by assumption, we have

$$\mathbb{E}[1_{(\mathcal{E}^c)}\|\theta_1 - \theta^*\|_2] \leq \mathbb{P}(\mathcal{E}^c)R \stackrel{(i)}{=} \mathcal{O}(Rn^{-k/2})$$

for any $k \leq \min\{k_2, k_1, k_0\}$, where step (i) follows from the inequality (27). Now, Hölder's inequality also yields that

$$\begin{aligned}\mathbb{E}[1_{(\mathcal{E}^c)}\|\Sigma^{-1}(P + Q)\Sigma^{-1}\nabla F_1(\theta^*)\|_2] &\leq \mathbb{E}[1_{(\mathcal{E}^c)}\|\Sigma^{-1}(P + Q)\|_2\|\nabla F_1(\theta^*)\|_2] \\ &\leq \sqrt{\mathbb{P}(\mathcal{E}^c)}\mathbb{E}\left[\|\Sigma^{-1}(P + Q)\|_2^4\right]^{1/4}\mathbb{E}\left[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^4\right]^{1/4}.\end{aligned}$$

Recalling Lemmas 2 and 4, we have $\mathbb{E}[\|\Sigma^{-1}(P + Q)\|_2^4] = \mathcal{O}(\log^2(d)n^{-2})$, and we similarly have $\mathbb{E}[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^4] = \mathcal{O}(n^{-2})$. Lastly, we have $\mathbb{P}(\mathcal{E}^c) = \mathcal{O}(n^{-k})$ for $k \leq \min\{k_0, k_1, k_2\}$, whence we find that for any such k ,

$$\mathbb{E}[1_{(\mathcal{E}^c)}\|\Sigma^{-1}(P + Q)\Sigma^{-1}\nabla F_1(\theta^*)\|_2] = \mathcal{O}\left(\sqrt{\log(d)}n^{-k/4-1}\right).$$

We can similarly apply Lemma 3 to the last remaining term in the inequality (31) to obtain that for any $k \leq \min\{k_2, k_1, k_0\}$,

$$\begin{aligned}\mathbb{E}[1_{(\mathcal{E}^c)}\|-\Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)[- \Sigma^{-1}\nabla F_1(\theta^*) + \Sigma^{-1}(P + Q)(\theta_1 - \theta^*)]\|_2] \\ = \mathcal{O}(n^{-k/2} + n^{-k/4-1}).\end{aligned}$$

Applying these two bounds, we find that

$$\|\mathbb{E}[\theta_1 - \theta^*]\|_2 \leq \|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}((P + Q)(w_1 - w^*) - \nabla F_1(\theta^*))]\|_2 + \mathcal{O}(n^{-k}) \quad (32)$$

for any k such that $k \leq \min\{k_0, k_1, k_2\}/2$ and $k \leq \min\{k_0, k_1, k_2\}/4 + 1$.

In the remainder of the proof, we show that part of the bound (32) still consists only of higher-order terms, leaving us with an expression not involving $\theta_1 - \theta^*$. To that end, note that

$$\mathbb{E}[\|\Sigma^{-1}(P + Q)\Sigma^{-1}(P + Q)(\theta_1 - \theta^*)\|_2^2] = \mathcal{O}(n^{-3})$$

by three applications of Hölder's inequality, the fact that $\|Ax\|_2 \leq \|A\|_2 \|x\|_2$, and Lemmas 3 and 4. Coupled with our bound (32), we use the fact that $(a + b)^2 \leq 2a^2 + 2b^2$ to obtain

$$\|\mathbb{E}[\theta_1 - \theta^*]\|_2^2 \leq 2 \|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}\nabla F_1(\theta^*)]\|_2^2 + \mathcal{O}(n^{-3}). \quad (33)$$

We focus on bounding the remaining expectation. We have the following series of inequalities:

$$\begin{aligned} \|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}\nabla F_1(\theta^*)]\|_2 &\stackrel{(i)}{\leq} \mathbb{E}[\|\Sigma^{-1}(P + Q)\|_2 \|\Sigma^{-1}\nabla F_1(\theta^*)\|_2] \\ &\stackrel{(ii)}{\leq} \left(\mathbb{E}[\|\Sigma^{-1}(P + Q)\|_2^2] \mathbb{E}[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^2] \right)^{\frac{1}{2}} \\ &\stackrel{(iii)}{\leq} \left(2\mathbb{E}[\|\Sigma^{-1}P\|_2^2 + \|\Sigma^{-1}Q\|_2^2] \mathbb{E}[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^2] \right)^{\frac{1}{2}}. \end{aligned}$$

Here step (i) follows from Jensen's inequality and the fact that $\|Ax\|_2 \leq \|A\|_2 \|x\|_2$; step (ii) uses the Cauchy-Schwarz inequality; and step (iii) follows from the fact that $(a + b)^2 \leq 2a^2 + 2b^2$. We have already bounded the first two terms in the product in our proofs; in particular, Lemma 2 guarantees that $\mathbb{E}[\|P\|_2^2] \leq CH \log d/n$, while

$$\mathbb{E}[\|Q\|_2^2] \leq \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n L(X_i)^4\right]^{\frac{1}{2}} \mathbb{E}[\|\theta_1 - \theta^*\|_2^4]^{\frac{1}{2}} \leq C \frac{L^2 G^2}{(1 - \rho)^2 \lambda^2} \cdot n^{-1}$$

for some numerical constant C (recall Lemma 4). Summarizing our bounds on $\|P\|_2$ and $\|Q\|_2$, we have

$$\begin{aligned} &\|\mathbb{E}[\Sigma^{-1}(P + Q)\Sigma^{-1}\nabla F_1(\theta^*)]\|_2^2 \\ &\leq 2 \|\Sigma^{-1}\|_2^2 \left(\frac{2H^2(\log d + 1)}{n} + 2C \frac{L^2 G^2}{(1 - \rho)^2 \lambda^2 n} + \mathcal{O}(n^{-2}) \right) \mathbb{E}[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^2]. \end{aligned} \quad (34)$$

From Assumption C we know that $\mathbb{E}[\|\nabla F_1(\theta^*)\|_2^2] \leq G^2/n$ and $\|\Sigma^{-1}\|_2 \leq 1/\lambda$, and hence we can further simplify the bound (34) to obtain

$$\|\mathbb{E}[\theta_1 - \theta^*]\|_2^2 \leq \frac{C}{\lambda^2} \left(\frac{H^2 \log d + L^2 G^2 / (1 - \rho)^2}{n} \right) \mathbb{E}[\|\Sigma^{-1}\nabla F_1(\theta^*)\|_2^2] + \mathcal{O}(n^{-3})$$

for some numerical constant C , where we have applied our earlier inequality (33). Noting that we may (without loss of generality) take $\rho < \frac{1}{2}$, then applying this inequality with the bound (28) on $\mathbb{E}[\|\theta_1 - \theta^*\|_2^2]$ we previously proved to our decomposition (21) completes the proof.

B Proof of Theorem 2

Our proof of Theorem 2 begins with a simple inequality that mimics our first inequality (21) in the proof of Theorem 1. Recall the definitions of the averaged vector $\bar{\theta}_1$ and bootstrap averaged vector $\bar{\theta}_2$. Let θ_1 denote the minimizer of the (an arbitrary) empirical risk F_1 , and θ_2 denote the minimizer of the resampled empirical risk F_2 (from the same samples as θ_1). Then we have

$$\mathbb{E} \left[\left\| \frac{\bar{\theta}_1 - r\bar{\theta}_2}{1-r} - \theta^* \right\|_2^2 \right] \leq \left\| \mathbb{E} \left[\frac{\theta_1 - r\theta_2}{1-r} - \theta^* \right] \right\|_2^2 + \frac{1}{m} \mathbb{E} \left[\left\| \frac{\theta_1 - r\theta_2}{1-r} - \theta^* \right\|_2^2 \right]. \quad (35)$$

Thus, parallel to our proof of Theorem 1, it suffices to bound the two terms in the decomposition (35) separately. Specifically, we prove the following two lemmas.

Lemma 5. *Under the conditions of Theorem 2,*

$$\left\| \mathbb{E} \left[\frac{\theta_1 - r\theta_2}{1-r} - \theta^* \right] \right\|_2^2 = \mathcal{O} \left(\frac{1}{r(1-r)^2} n^{-3} \right). \quad (36)$$

Lemma 6. *Under the conditions of Theorem 2,*

$$\mathbb{E} \left[\left\| \theta_1 - \theta^* - r(\theta_2 - \theta^*) \right\|_2^2 \right] \leq (2+3r) \mathbb{E} \left[\left\| \nabla^2 F_0(\theta^*)^{-1} \nabla F_1(\theta^*) \right\|_2^2 \right] + \mathcal{O}(n^{-2}) \quad (37)$$

In conjunction, Lemmas 5 and 6 coupled with the decomposition (35) yield the desired claim. Indeed, applying each of the lemmas to the decomposition (35), we see that

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{\bar{\theta}_1 - r\bar{\theta}_2}{1-r} - \theta^* \right\|_2^2 \right] &\leq \frac{2+3r}{(1-r)^2 m} \mathbb{E} \left[\left\| \nabla^2 F_0(\theta^*)^{-1} \nabla F_1(\theta^*) \right\|_2^2 \right] \\ &\quad + \mathcal{O} \left(\frac{1}{(1-r)^2} m^{-1} n^{-2} \right) + \mathcal{O} \left(\frac{1}{r(1-r)^2} n^{-3} \right), \end{aligned}$$

which is the statement of Theorem 2.

Lemma 6 is quite similar to the results used in the proof of Theorem 1. Lemma 5 is more subtle and key to our argument, as it shows that by including the rate- r -subsampled estimates θ_2 and correcting using the bootstrap-estimated bias $r(\theta_1 - \theta_2)$, we can obtain an estimate of θ^* which has expectation of smaller order than was possible in Theorem 1.

With this set-up, our proof now proceeds in three steps. First, we give asymptotic expansions of the differences $\theta_1 - \theta^*$ and $\theta_2 - \theta^*$. These expansions allow us to show a key result on the bias-correction provided by the bootstrap, since the expansions $\theta_2 - \theta^*$ and $\theta_1 - \theta^*$ turn out to have only second-order differences, which yields Lemma 5. We prove Lemma 6 using techniques analogous to those for Theorem 1, but we defer its (somewhat more technical) proof to Appendix G.

Throughout the rest of the proof, we use the notation

$$Y = Y' + \mathcal{R}_k$$

for some random variables Y and Y' to mean that there exists a random variable Z such that $Y = Y' + Z$ and $\mathbb{E}[\|Z\|_2^2] = \mathcal{O}(n^{-k})$.¹ The symbol \mathcal{R}_k may indicate different random variables

¹ Formally, in our proof this will mean that there exist random vectors Y , Y' , and Z that are measurable with respect to the σ -field $\sigma(X_1, \dots, X_n)$, where $Y = Y' + Z$ and $\mathbb{E}[\|Z\|_2^2] = \mathcal{O}(n^{-k})$.

throughout a proof and is notational shorthand for a moment-based big-O notation. We also remark that if we have $\mathbb{E}[\|Z\|_2^2] = \mathcal{O}(a^k n^{-k})$, we have $Z = a^{k/2} \mathcal{R}_k$, since $(a^{k/2})^2 = a^k$. For shorthand, we also say that $\mathbb{E}[Z] = \mathcal{O}(h(n))$ if $\|\mathbb{E}[Z]\|_2 = \mathcal{O}(h(n))$, which implies that if $Z = \mathcal{R}_k$ then $\mathbb{E}[Z] = \mathcal{O}(n^{-k/2})$, since

$$\|\mathbb{E}[Z]\|_2 \leq \sqrt{\mathbb{E}[\|Z\|_2^2]} = \mathcal{O}(n^{-k/2}).$$

B.1 Optimization Error Expansion

In this section, we derive a sharper asymptotic expansion of the optimization errors $\theta_1 - \theta^*$. Recall our definition of the Kronecker product \otimes , where for vectors u, v we have $u \otimes v = uv^\top$. With this notation, we have the following expansion of $\theta_1 - \theta^*$.

Lemma 7. *Under the conditions of Theorem 2, we have*

$$\begin{aligned} \theta_1 - \theta^* &= -\Sigma^{-1} \nabla F_1(\theta^*) + \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*) \\ &\quad - \Sigma^{-1} \nabla^3 F_0(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) + \mathcal{R}_3. \end{aligned} \quad (38)$$

We provide a proof of Lemma 7 in Appendix H. The lemma requires somewhat more careful moment control over the expansion $\theta_1 - \theta^*$, which leads to some technical difficulty, but is similar in spirit to the results used to prove Theorem 1.

An immediately analogous result to Lemma 7 follows for our sub-sampled bootstrap estimators. Since we use rn samples to compute θ_2 , the second level estimator, we find

Lemma 8. *Under the conditions of Theorem 2, we have*

$$\begin{aligned} \theta_2 - \theta^* &= -\Sigma^{-1} \nabla F_2(\theta^*) + \Sigma^{-1} (\nabla^2 F_2(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_2(\theta^*) \\ &\quad - \Sigma^{-1} \nabla^3 F_0(\theta^*) ((\Sigma^{-1} \nabla F_2(\theta^*)) \otimes (\Sigma^{-1} \nabla F_2(\theta^*))) + r^{-3/2} \mathcal{R}_3. \end{aligned}$$

B.2 Bias Correction

Now that we have given Taylor expansions that describe the behavior of $\theta_1 - \theta^*$ and $\theta_2 - \theta^*$, we can prove Lemmas 5 and 6 (though, as noted earlier, we defer the proof of Lemma 6 to Appendix G). The key insight is that expectations of terms involving $\nabla F_2(\theta^*)$ are nearly the same as expectations of terms involving $\nabla F_1(\theta^*)$, except that some corrections for the sampling ratio r are necessary.

We begin by noting that

$$\frac{\theta_1 - r\theta_2}{1-r} - \theta^* = \frac{\theta_1 - \theta^*}{1-r} - r \frac{\theta_2 - \theta^*}{1-r}. \quad (39)$$

In Lemmas 7 and 8, we derived expansions for each of the right hand side terms, and since

$$\mathbb{E}[\Sigma^{-1} \nabla F_1(\theta^*)] = 0 \quad \text{and} \quad \mathbb{E}[\Sigma^{-1} \nabla F_2(\theta^*)] = 0,$$

Lemmas 7 and 8 coupled with the rewritten bootstrap correction (39) yield

$$\begin{aligned} \mathbb{E}[\theta_1 - \theta^* - r(\theta_2 - \theta^*)] &= r \mathbb{E}[\Sigma^{-1} (\nabla^2 F_2(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_2(\theta^*)] \\ &\quad - \mathbb{E}[\Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*)] \\ &\quad + r \mathbb{E}[\Sigma^{-1} \nabla^3 F_0(\theta^*) ((\Sigma^{-1} \nabla F_2(\theta^*)) \otimes (\Sigma^{-1} \nabla F_2(\theta^*)))] \\ &\quad - \mathbb{E}[\Sigma^{-1} \nabla^3 F_0(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*)))] \\ &\quad + r^{-1/2} \mathcal{O}(n^{-3/2}). \end{aligned} \quad (40)$$

Here the remainder terms follow because of the $r^{-3/2}\mathcal{R}_3$ term on $\theta_2 - \theta^*$. We can now give our proof of Lemma 5.

Proof of Lemma 5 To prove the claim in the lemma, it suffices to show that

$$r\mathbb{E}[\Sigma^{-1}(\nabla^2 F_2(\theta^*) - \Sigma)\Sigma^{-1}\nabla F_2(\theta^*)] = \mathbb{E}[\Sigma^{-1}(\nabla^2 F_1(\theta^*) - \Sigma)\Sigma^{-1}\nabla F_1(\theta^*)] \quad (41)$$

and

$$\begin{aligned} r\mathbb{E}[\Sigma^{-1}\nabla^3 F_0(\theta^*) ((\Sigma^{-1}\nabla F_2(\theta^*)) \otimes (\Sigma^{-1}\nabla F_2(\theta^*)))] \\ = \mathbb{E}[\Sigma^{-1}\nabla^3 F_0(\theta^*) ((\Sigma^{-1}\nabla F_1(\theta^*)) \otimes (\Sigma^{-1}\nabla F_1(\theta^*)))] \end{aligned} \quad (42)$$

Indeed, these two claims combined with the expansion (40) yield the bound (36) in Lemma 5 immediately.

We first consider the difference (41). To make things notationally simpler, we define functions let $A : \mathcal{X} \rightarrow \mathbb{R}^{d \times d}$ and $B : \mathcal{X} \rightarrow \mathbb{R}^d$ via $A(x) := \Sigma^{-1}(\nabla^2 f(\theta^*; x) - \Sigma)$ and $B(x) := \Sigma^{-1}\nabla f(\theta^*; x)$. If we let $S_1 = \{X_1, \dots, X_n\}$ be the original (non-bootstrap) samples and $S_2 = \{Y_1, \dots, Y_{rn}\}$ be the subsampled dataset, we must show

$$r\mathbb{E}\left[\frac{1}{(rn)^2} \sum_{i,j}^{rn} A(Y_i)B(Y_j)\right] = \mathbb{E}\left[\frac{1}{n^2} \sum_{i,j}^n A(X_i)B(X_j)\right].$$

But since the Y_i are sampled without replacement, and $\mathbb{E}[A(X_i)] = 0$ and $\mathbb{E}[B(X_i)] = 0$, we find that $\mathbb{E}[A(Y_i)B(Y_j)] = 0$ for $i \neq j$, and thus

$$\sum_{i,j}^{rn} \mathbb{E}[A(Y_i)B(Y_j)] = \sum_{i=1}^{rn} \mathbb{E}[A(Y_i)B(Y_i)] = rn\mathbb{E}[A(Y_1)B(Y_1)].$$

In particular, we see that the equality (41) holds:

$$\frac{r}{(rn)^2} \sum_{i,j}^{rn} \mathbb{E}[A(Y_i)B(Y_j)] = \frac{r}{rn} \mathbb{E}[A(Y_1)B(Y_1)] = \frac{1}{n} \mathbb{E}[A(X_1)B(X_1)] = \frac{1}{n^2} \sum_{i,j}^n \mathbb{E}[A(X_i)B(X_j)].$$

The statement (42) follows from completely parallel arguments; we omit details for brevity. \square

C Proof of Theorem 3

We begin by recalling that if θ^n denotes the output of performing stochastic gradient on one machine, then from the inequality (21) we have the upper bound

$$\mathbb{E}[\|\bar{\theta}^n - \theta^*\|_2^2] \leq \frac{1}{m} \mathbb{E}[\|\theta^n - \theta^*\|_2^2] + \|\mathbb{E}[\theta^n - \theta^*]\|_2^2.$$

To prove the error bound (18), it thus suffices to prove the inequalities

$$\mathbb{E}[\|\theta^n - \theta^*\|_2^2] \leq \frac{\alpha G^2}{\lambda^2 n}, \quad \text{and} \quad (43a)$$

$$\|\mathbb{E}[\theta^n - \theta^*]\|_2^2 \leq \frac{\beta^2}{n^{3/2}}. \quad (43b)$$

Before proving the theorem, we introduce some notation and a few preliminary results. Let $g_t = \nabla f(\theta^t; X_t)$ be the gradient of the t^{th} sample in stochastic gradient descent, where we consider running SGD on a single machine. We also let

$$\Pi(v) := \operatorname{argmin}_{\theta \in \Theta} \left\{ \|\theta - v\|_2^2 \right\}$$

denote the projection of the point v onto the domain Θ .

We now state a lemma due to Rakhlin et al. [24], which gives sharp rates on the convergence of the iterates $\{\theta^t\}$ in stochastic gradient descent.

Lemma 9 (Rakhlin, Shamir, Sridharan). *Assume that $\mathbb{E}[\|g_t\|_2^2] \leq G^2$ for all t . Choosing $\eta_t = \frac{c}{\lambda t}$ for some $c > 1/2$, for any $t \in \mathbb{N}$ we have*

$$\mathbb{E} \left[\|\theta^t - \theta^*\|_2^2 \right] \leq \frac{\alpha G^2}{\lambda^2 t} \quad \text{where} \quad \alpha = \max(4, c/(2 - 1/c)).$$

With these ingredients, we can now turn to the proof of Theorem 3. Lemma 9 gives the inequality (43a), so it remains to prove that $\bar{\theta}^n$ has the smaller bound (43b) on its bias. To that end, recall the neighborhood $U_\rho \subset \Theta$ in Assumption E, and note that

$$\begin{aligned} \theta^{t+1} - \theta^* &= \Pi(\theta^t - \eta_t g_t - \theta^*) \\ &= \theta^t - \eta_t g_t - \theta^* + 1_{(\theta^{t+1} \notin U_\rho)} (\Pi(\theta^t - \eta_t g_t) - (\theta^t - \eta_t g_t)) \end{aligned}$$

since when $\theta \in U_\rho$, we have $\Pi(\theta) = \theta$. Consequently, an application of the triangle inequality gives

$$\|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 \leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \mathbb{E}[\|(\Pi(\theta^t - \eta_t g_t) - (\theta^t - \eta_t g_t))1_{(\theta^{t+1} \notin U_\rho)}\|_2].$$

By the definition of the projection and the fact that $\theta^t \in \Theta$, we additionally have

$$\|\Pi(\theta^t - \eta_t g_t) - (\theta^t - \eta_t g_t)\|_2 \leq \|\theta^t - (\theta^t - \eta_t g_t)\|_2 \leq \eta_t \|g_t\|_2.$$

Thus, by applying Hölder's inequality (with the conjugate choices $(p, q) = (4, \frac{4}{3})$) and Assumption E, we have

$$\begin{aligned} \|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 &\leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \eta_t \mathbb{E}[\|g_t\|_2 1_{(\theta^{t+1} \notin U_\rho)}] \\ &\leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \eta_t \sqrt[4]{\mathbb{E}[\|g_t\|_2^4]} \left(\mathbb{E}[1_{(\theta^t \notin U_\rho)}^{4/3}] \right)^{3/4} \\ &\leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \eta_t G (\mathbb{P}(\theta^t \notin U_\rho))^{3/4} \\ &\leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \eta_t G \left(\frac{\mathbb{E}[\|\theta^{t+1} - \theta^*\|_2^2]}{\rho^2} \right)^{3/4}, \end{aligned} \tag{44}$$

the inequality (44) following from an application of Markov's inequality. By applying Lemma 9, we finally obtain

$$\begin{aligned} \|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 &\leq \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \eta_t G \left(\frac{\alpha G^2}{\lambda^2 \rho^2 t} \right)^{3/4} \\ &= \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 + \frac{c \alpha^{3/4} G^{5/2}}{\lambda^{5/2} \rho^{3/2}} \cdot \frac{1}{t^{7/4}}. \end{aligned} \tag{45}$$

Now we turn to controlling the rate at which $\theta^t - \eta_t g_t$ goes to zero. Let $f_t(\cdot) = f(\cdot; X_t)$ be shorthand for the loss evaluated on the t^{th} data point. By defining $r_t = g_t - \nabla f_t(\theta^*) - \nabla^2 f_t(\theta^*)(\theta^t - \theta^*)$, a bit of algebra yields

$$g_t = \nabla f_t(\theta^*) + \nabla^2 f_t(\theta^*)(\theta^t - \theta^*) + r_t.$$

Since θ^t belongs to the σ -field of X_1, \dots, X_{t-1} , the Hessian $\nabla^2 f_t(\theta^*)$ is (conditionally) independent of θ^t and

$$\mathbb{E}[g_t] = \nabla^2 F_0(\theta^*)\mathbb{E}[\theta^t - \theta^*] + \mathbb{E}[r_t 1_{(\theta^t \in U_\rho)}] + \mathbb{E}[r_t 1_{(\theta^t \notin U_\rho)}]. \quad (46)$$

If $\theta^t \in U_\rho$, then Taylor's theorem implies that r_t is the Lagrange remainder

$$r_t = (\nabla^2 f_t(\theta') - \nabla^2 f_t(\theta^*))(\theta' - \theta^*),$$

where $\theta' = \kappa\theta^t + (1 - \kappa)\theta^*$ for some $\kappa \in [0, 1]$. Applying Assumption E and Hölder's inequality, we find that since θ^t is conditionally independent of X_t ,

$$\begin{aligned} \mathbb{E} \left[\|r_t 1_{(\theta^t \in U_\rho)}\|_2 \right] &\leq \mathbb{E} \left[\|\nabla^2 f(\theta'; X_t) - \nabla^2 f(\theta^*; X_t)\| \|\theta^t - \theta^*\|_2 1_{(\theta^t \in U_\rho)} \right] \\ &\leq \mathbb{E} \left[L(X_t) \|\theta^t - \theta^*\|_2^2 \right] = \mathbb{E}[L(X_t)] \mathbb{E}[\|\theta^t - \theta^*\|_2^2] \\ &\leq L \mathbb{E} \left[\|\theta^t - \theta^*\|_2^2 \right] \leq \frac{\alpha L G^2}{\lambda^2 t}. \end{aligned}$$

On the other hand, when $\theta^t \notin U_\rho$, we have the following sequence of inequalities:

$$\begin{aligned} \mathbb{E} \left[\|r_t 1_{(\theta^t \notin U_\rho)}\|_2 \right] &\stackrel{(i)}{\leq} \sqrt[4]{\mathbb{E}[\|r_t\|_2^4]} (\mathbb{P}(\theta^t \notin U_\rho))^{3/4} \\ &\stackrel{(ii)}{\leq} \sqrt[4]{3^3 \left(\mathbb{E}[\|g_t\|_2^4] + \mathbb{E}[\|\nabla f_t(\theta^*)\|_2^4] + \mathbb{E}[\|\nabla^2 f_t(\theta^*)(\theta^t - \theta^*)\|_2^4] \right)} (\mathbb{P}(\theta^t \notin U_\rho))^{3/4} \\ &\leq 3^{3/4} \sqrt[4]{G^4 + G^4 + H^4 R^4} (\mathbb{P}(\theta^t \notin U_\rho))^{3/4} \\ &\stackrel{(iii)}{\leq} 3(G + HR) \left(\frac{\alpha G^2}{\lambda^2 \rho^2 t} \right)^{3/4}. \end{aligned}$$

Here step (i) follows from Hölder's inequality (again applied with the conjugates $(p, q) = (4, \frac{4}{3})$); step (ii) follows from Jensen's inequality, since $(a + b + c)^4 \leq 3^3(a^4 + b^4 + c^4)$; and step (iii) follows from Markov's inequality, as in the bounds (44) and (45). Combining our two bounds on r_t , we find that

$$\mathbb{E}[\|r_t\|_2] \leq \frac{\alpha L G^2}{\lambda^2 t} + \frac{3\alpha^{3/4} G^{3/2} (G + HR)}{\lambda^{3/2} \rho^{3/2}} \cdot \frac{1}{t^{3/4}}. \quad (47)$$

By combining the expansion (46) with the bound (47), we find that

$$\begin{aligned} \|\mathbb{E}[\theta^t - \eta_t g_t - \theta^*]\|_2 &= \|\mathbb{E}[(I - \eta_t \nabla^2 F_0(\theta^*))(\theta^t - \theta^*) + \eta_t r_t]\|_2 \\ &\leq \|\mathbb{E}[(I - \eta_t \nabla^2 F_0(\theta^*))(\theta^t - \theta^*)]\|_2 + \frac{c\alpha L G^2}{\lambda^3 t^2} + \frac{3c\alpha^{3/4} G^{3/2} (G + HR)}{\lambda^{5/2} \rho^{3/2}} \cdot \frac{1}{t^{7/4}}. \end{aligned}$$

Using the earlier bound (45), this inequality then yields

$$\|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 \leq \|I - \eta_t \nabla^2 F_0(\theta^*)\|_2 \|\mathbb{E}[\theta^t - \theta^*]\|_2 + \frac{c\alpha^{3/4} G^{3/2}}{\lambda^{5/2} t^{7/4}} \left(\frac{\alpha^{1/4} L G^{1/2}}{\lambda^{1/2} t^{1/4}} + \frac{4G + HR}{\rho^{3/2}} \right).$$

We now complete the proof via an inductive argument using our immediately preceding bounds. Our reasoning follows a similar induction given by Rakhlin et al. [24]. First, note that by strong convexity and our condition that $\|\nabla^2 F_0(\theta^*)\| \leq H$, we have

$$\|I - \eta_t \nabla^2 F_0(\theta^*)\| = 1 - \eta_t \lambda_{\min}(\nabla^2 F_0(\theta^*)) \leq 1 - \eta_t \lambda$$

whenever $1 - \eta_t H \geq 0$. Define $\tau_0 = \lceil cH/\lambda \rceil$; then for $t \geq t_0$ we obtain

$$\|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 \leq (1 - c/t) \|\mathbb{E}[\theta^t - \theta^*]\|_2 + \frac{1}{t^{7/4}} \cdot \frac{c\alpha^{3/4}G^{3/2}}{\lambda^{5/2}} \left(\frac{\alpha^{1/4}LG^{1/2}}{\lambda^{1/2}t^{1/4}} + \frac{4G + HR}{\rho^{3/2}} \right). \quad (48)$$

For shorthand, we define two intermediate variables

$$a_t = \|\mathbb{E}(\theta^t - \theta^*)\|_2 \quad \text{and} \quad b = \frac{c\alpha^{3/4}G^{3/2}}{\lambda^{5/2}} \left(\frac{\alpha^{1/4}LG^{1/2}}{\lambda^{1/2}} + \frac{4G + HR}{\rho^{3/2}} \right).$$

Inequality (48) then implies the inductive relation $a_{t+1} \leq (1 - c/t)a_t + b/t^{7/4}$. Now we show that by defining $\beta = \max\{\tau_0 R, b/(c-1)\}$, we have $a_t \leq \beta/t^{3/4}$. Indeed, it is clear that $a_1 \leq \tau_0 R$. Using the inductive hypothesis, we then have

$$a_{t+1} \leq \frac{(1 - c/t)\beta}{t^{3/4}} + \frac{b}{t^{7/4}} = \frac{\beta(t-1)}{t^{7/4}} - \frac{\beta(c-1) - b}{t^2} \leq \frac{\beta(t-1)}{t^{7/4}} \leq \frac{\beta}{(t+1)^{3/4}}.$$

This completes the proof of the inequality (43b). \square

Remark If we assume k th moment bounds instead of 4th, i.e. $\mathbb{E}[\|\nabla^2 f(\theta^*; X)\|_2^k] \leq H^k$ and $\mathbb{E}[\|g_t\|_2^k] \leq G^k$, we find the following analogue of the bound (48):

$$\begin{aligned} \|\mathbb{E}[\theta^{t+1} - \theta^*]\|_2 &\leq (1 - c/t) \|\mathbb{E}[\theta^t - \theta^*]\|_2 \\ &\quad + \frac{1}{t^{\frac{2k-1}{k}}} \cdot \frac{c\alpha^{\frac{k-1}{k}}G^{\frac{2k-2}{k}}}{\lambda^{\frac{3k-2}{k}}} \left[\frac{(54^{1/k} + 1)G + 54^{1/k}HR}{\rho^{\frac{2k-2}{k}}} + \frac{\alpha^{1/k}LG^{2/k}}{\lambda^{2/k}t^{1/k}} \right]. \end{aligned}$$

In this case, if we define

$$b = \frac{c\alpha^{\frac{k-1}{k}}G^{\frac{2k-2}{k}}}{\lambda^{\frac{3k-2}{k}}} \left[\frac{(54^{1/k} + 1)G + 54^{1/k}HR}{\rho^{\frac{2k-2}{k}}} + \frac{\alpha^{1/k}LG^{2/k}}{\lambda^{2/k}} \right] \quad \text{and} \quad \beta = \max \left\{ \tau_0 R, \frac{b}{c-1} \right\},$$

we have the same result except we obtain the bound $\|\mathbb{E}[\theta^n - \theta^*]\|_2^2 \leq \beta^2/n^{\frac{2k-2}{k}}$.

D The necessity of smoothness

Here we show that some version of the smoothness conditions presented in Assumption C are necessary for averaging methods to attain better mean-squared error than using only the n samples on a single processor. We use the loss function (10), and set $n_0 = \sum_{i=1}^n 1_{(X_i=0)}$ to be the count of 0 samples. Using θ_1 as shorthand for $\theta_{1,i}$, we see by inspection that the empirical minimizer θ_1 is

$$\theta_1 = \begin{cases} \frac{n_0}{n} - \frac{1}{2} & \text{when } n_0 \leq n/2 \\ 1 - \frac{n}{2n_0} & \text{otherwise.} \end{cases}$$

For simplicity, we may assume that n is odd. In this case, we obtain that

$$\begin{aligned}\mathbb{E}[\theta_1] &= \frac{1}{4} + \mathbb{E}\left[\frac{n_0}{n}1_{(n_0 < n/2)}\right] - \mathbb{E}\left[\frac{n}{2n_0}1_{(n_0 > n/2)}\right] \\ &= \frac{1}{4} + \frac{1}{2^n} \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{i} \frac{i}{n} - \frac{1}{2^n} \sum_{i=\lfloor n/2 \rfloor}^n \binom{n}{i} \frac{n}{2i} = \frac{1}{4} + \frac{1}{2^n} \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{i} \left[\frac{i}{n} - \frac{n}{2(n-i)} \right]\end{aligned}$$

by the symmetry of the binomial. Adding and subtracting $\frac{1}{2}$ from the term within the braces, noting that $P(n_0 < n/2) = 1/2$, we have the equality

$$\mathbb{E}[\theta_1] = \frac{1}{2^n} \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{i} \left[\frac{i}{n} - \frac{n}{2(n-i)} + \frac{1}{2} \right] = \frac{1}{2^n} \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{i} \frac{i(n-2i)}{2n(n-i)}.$$

If Z is distributed normally with mean $1/2$ and variance $1/n$, then an asymptotic expansion of the binomial distribution yields

$$\begin{aligned}\left(\frac{1}{2}\right)^n \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{i} \frac{i(n-2i)}{2n(n-i)} &= \mathbb{E}\left[\frac{Z(1-2Z)}{2-2Z} \mid 0 \leq Z \leq \frac{1}{2}\right] + o(n^{-1/2}) \\ &\geq \frac{1}{2} \mathbb{E}\left[Z - 2Z^2 \mid 0 \leq Z \leq \frac{1}{2}\right] + o(n^{-1/2}) = \Omega(n^{-\frac{1}{2}}),\end{aligned}$$

the final equality following from standard calculations.

E Proof of Lemma 1

We first prove that under the conditions given in the lemma statement, the function F_1 is $(1-\rho)\lambda$ -strongly convex over the ball $U := \{\theta \in \mathbb{R}^d : \|\theta - \theta^*\|_2 < \delta_\rho\}$ around θ^* . Indeed, fix $\gamma \in U$, then use the triangle inequality to conclude that

$$\begin{aligned}\|\nabla^2 F_1(\gamma) - \nabla^2 F_0(\theta^*)\|_2 &\leq \|\nabla^2 F_1(\gamma) - \nabla^2 F_1(\theta^*)\|_2 + \|\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*)\|_2 \\ &\leq L \|\gamma - \theta^*\|_2 + \frac{\rho\lambda}{2}.\end{aligned}$$

Here we used Assumption C on the first term and the fact that the event \mathcal{E}_1 holds on the second. By our choice of $\delta_\rho \leq \rho\lambda/4L$, this final term is bounded by $\lambda\rho$. In particular, we have

$$\nabla^2 F_0(\theta^*) \succeq \lambda I \quad \text{so} \quad \nabla^2 F_0(\gamma) \succeq \lambda I - \rho\lambda I = (1-\rho)\lambda I,$$

which proves that F_1 is $(1-\rho)\lambda$ -strongly convex on the ball U .

In order to prove the conclusion of the lemma, we argue that since F_1 is (locally) strongly convex, if the function F_1 has small gradient at the point θ^* , it must be the case that the minimizer θ_1 of F_1 is near θ^* . Then we can employ reasoning similar to standard analyses of optimality for globally strongly convex functions (e.g. [3], Chapter 9). By definition of (the local) strong convexity on the set U , for any $\theta' \in \Theta$, we have

$$F_1(\theta') \geq F_1(\theta^*) + \langle \nabla F_1(\theta^*), \theta' - \theta^* \rangle + \frac{(1-\rho)\lambda}{2} \min\left\{\|\theta^* - \theta'\|_2^2, \delta_\rho^2\right\}.$$

Rewriting this inequality, we find that

$$\begin{aligned} \min \left\{ \|\theta^* - \theta'\|_2^2, \delta_\rho^2 \right\} &\leq \frac{2}{(1-\rho)\lambda} [F_1(\theta') - F_1(\theta^*) + \langle \nabla F_1(\theta^*), \theta' - \theta^* \rangle] \\ &\leq \frac{2}{(1-\rho)\lambda} [F_1(\theta') - F_1(\theta^*) + \|\nabla F_1(\theta^*)\|_2 \|\theta' - \theta^*\|_2]. \end{aligned}$$

Dividing each side by $\|\theta' - \theta^*\|_2$, then noting that we may set $\theta' = \kappa\theta_1 + (1-\kappa)\theta^*$ for any $\kappa \in [0, 1]$, we have

$$\min \left\{ \kappa \|\theta_1 - \theta^*\|_2, \frac{\delta_\rho^2}{\kappa \|\theta_1 - \theta^*\|_2} \right\} \leq \frac{2[F_1(\kappa\theta_1 + (1-\kappa)\theta^*) - F_1(\theta^*)]}{\kappa(1-\rho)\lambda \|\theta_1 - \theta^*\|_2} + \frac{2\|\nabla F_1(\theta^*)\|_2}{(1-\rho)\lambda}.$$

Of course, $F_1(\theta_1) < F_1(\theta^*)$ by assumption, so we find that for any $\kappa \in (0, 1)$ we have the strict inequality

$$\min \left\{ \kappa \|\theta_1 - \theta^*\|_2, \frac{\delta_\rho^2}{\kappa \|\theta_1 - \theta^*\|_2} \right\} < \frac{2\|\nabla F_1(\theta^*)\|_2}{(1-\rho)\lambda} \leq \delta_\rho,$$

the last inequality following from the definition of \mathcal{E}_2 . Since this holds for any $\kappa \in (0, 1)$, if $\|\theta_1 - \theta^*\|_2 > \delta_\rho$, we may set $\kappa = \delta_\rho / \|\theta_1 - \theta^*\|_2$, which would yield a contradiction. Thus, we have $\|\theta_1 - \theta^*\|_2 \leq \delta_\rho$, and by our earlier inequalities,

$$\|\theta_1 - \theta^*\|_2^2 \leq \frac{2}{(1-\rho)\lambda} [F_1(\theta_1) - F_1(\theta^*) + \|\nabla F_1(\theta^*)\|_2 \|\theta_1 - \theta^*\|_2] \leq \frac{2\|\nabla F_1(\theta^*)\|_2}{(1-\rho)\lambda} \|\theta_1 - \theta^*\|_2.$$

Dividing by $\|\theta_1 - \theta^*\|_2$ completes the proof. \square

F Moment bounds

In this appendix, we state two useful moment bounds, showing how they combine to provide a proof of Lemma 2. The two lemmas are a vector and a non-commutative matrix variant of the classical Rosenthal inequalities. We have

Lemma 10 (Chen et al. [4], Theorem A.1(2)). *Let $X_i \in \mathbb{R}^{d \times d}$ be independent symmetrically distributed Hermitian matrices. Then*

$$\mathbb{E} \left[\left\| \sum_{i=1}^n X_i \right\|^k \right]^{1/k} \leq \sqrt{2e \log d} \left\| \left(\sum_{i=1}^n \mathbb{E} [X_i^2] \right)^{1/2} \right\| + 2e \log d \left(\mathbb{E} [\max_i \|X_i\|^k] \right)^{1/k}.$$

Lemma 11 (de Acosta [5], Theorem 2.1). *Let $k \geq 2$ and X_i be a sequence of independent random vectors in a separable Banach space with norm $\|\cdot\|$ and $\mathbb{E}[\|X_i\|^k] < \infty$. There exists a finite constant C_k such that*

$$\mathbb{E} \left[\left\| \sum_{i=1}^n X_i \right\|^k \right] - \mathbb{E} \left[\left\| \sum_{i=1}^n X_i \right\|^2 \right]^k \leq C_k \left[\left(\sum_{i=1}^n \mathbb{E} [\|X_i\|^2] \right)^{k/2} + \sum_{i=1}^n \mathbb{E} [\|X_i\|^k] \right].$$

Now we turn to our proof Lemma 2. To prove the first bound (26a), let $2 \leq k \leq k_0$ and note that by Jensen's inequality, we have

$$\mathbb{E}[\|\nabla F_1(\theta^*)\|_2^k] \leq 2^{k-1} \mathbb{E} \left[|\|\nabla F_1(\theta^*)\|_2 - \mathbb{E}[\|\nabla F_1(\theta^*)\|_2]|^k \right] + 2^{k-1} \mathbb{E} [\|\nabla F_1(\theta^*)\|_2^k].$$

Again applying Jensen's inequality, $\mathbb{E}[\|\nabla f(\theta^*; X)\|_2^2] \leq G^2$. Thus by recalling the definition $\nabla F_1(\theta^*) = \frac{1}{n} \sum_{i=1}^n \nabla f(\theta^*; X_i)$ and applying the inequality $\mathbb{E}[\|\nabla F_1(\theta^*)\|_2] \leq \mathbb{E}[\|\nabla F_1(\theta^*)\|_2^2]^{1/2} \leq n^{-1/2}G$, we see that Lemma 11 implies

$$\begin{aligned} & \mathbb{E} [\|\nabla F_1(\theta^*)\|_2^k] \\ & \leq 2^{k-1} C_k \left[\left(\frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^2] \right)^{k/2} + \frac{1}{n^k} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^k] \right] + 2^{k-1} \mathbb{E}[\|\nabla F_1(\theta^*)\|_2^k] \\ & \leq 2^{k-1} C_k n^{-k/2} \left[\left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^2] \right)^{k/2} + n^{-k/2} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^k] \right] + 2^{k-1} n^{-k/2} G^k. \end{aligned}$$

Applying Jensen's inequality again to note that

$$\left(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^2] \right)^{k/2} \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla f(\theta^*; X_i)\|_2^{2k/2}]^{k/2} \leq G^k$$

completes the proof of the inequality (26a).

The proof of the bound (26b) requires a very slightly more delicate argument, involving symmetrization step. Define matrices $Z_i = \frac{1}{n} (\nabla^2 f(\theta^*; X_i) - \nabla^2 F_0(\theta^*))$. If $\varepsilon_i \in \{\pm 1\}$ are i.i.d. Rademacher variables independent of Z_i , then for any integer k in the interval $[2, k_2]$, a standard symmetrization argument (e.g. [14, Lemma 6.3]) implies that

$$\mathbb{E} \left[\left\| \sum_{i=1}^n Z_i \right\|^k \right]^{1/k} \leq 2 \mathbb{E} \left[\left\| \sum_{i=1}^n \varepsilon_i Z_i \right\|^k \right]^{1/k}. \quad (49)$$

Now we may apply Lemma 11, since the matrices $\varepsilon_i Z_i$ are Hermitian and symmetrically distributed; by expanding the definition of the Z_i , we find that

$$\begin{aligned} \mathbb{E} \left[\left\| \nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*) \right\|^k \right]^{1/k} & \leq 2 \sqrt{2e \log d} \left\| \left(\frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[(\nabla^2 f(\theta^*; X_i) - \nabla^2 F_0(\theta^*))^2] \right)^{1/2} \right\| \\ & \quad + 4e \log d \left(n^{-k} \mathbb{E}[\max_i \|\nabla^2 f(\theta^*; X_i) - \nabla^2 F_0(\theta^*)\|^k] \right)^{1/k}. \end{aligned}$$

Since the X_i are i.i.d., we have

$$\begin{aligned} \left\| \left(\frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[(\nabla^2 f(\theta^*; X_i) - \nabla^2 F_0(\theta^*))^2] \right)^{1/2} \right\| & = \left\| n^{-1/2} \mathbb{E} [(\nabla^2 f(\theta^*; X) - \nabla^2 F_0(\theta^*))^2]^{1/2} \right\| \\ & \leq n^{-1/2} \mathbb{E} \left[\left\| \nabla^2 f(\theta^*; X) - \nabla^2 F_0(\theta^*) \right\|^2 \right]^{1/2} \end{aligned}$$

by Jensen's inequality, since $\|A^{1/2}\| = \|A\|^{1/2}$ for positive semidefinite A . Finally, noting that

$$n^{-k} \mathbb{E} \left[\max_i \|\nabla^2 f(\theta^*; X_i) - \nabla^2 F_0(\theta^*)\|^k \right] \leq n^{-k} n \mathbb{E} \left[\|\nabla^2 f(\theta^*; X) - \nabla^2 F_0(\theta^*)\|^k \right] \leq n^{1-k} H^k$$

completes the proof of the second bound (26b).

G Proof of Lemma 6

The proof of Lemma 6 follows from that of Lemmas 7 and 8. We first claim that

$$\theta_1 - \theta^* = -\Sigma^{-1} \nabla F_1(\theta^*) + \mathcal{R}_2 \quad \text{and} \quad \theta_2 - \theta^* = -\Sigma^{-1} \nabla F_2(\theta^*) + r^{-1} \mathcal{R}_2. \quad (50)$$

The proofs of both claims similar, so we focus on proving the second statement. Using the inequality $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ and Lemma 8, we see that

$$\begin{aligned} \mathbb{E} \left[\|\theta_2 - \theta^* + \Sigma^{-1} \nabla F_2(\theta^*)\|_2^2 \right] &\leq 3 \mathbb{E} \left[\|\Sigma^{-1} (\nabla^2 F_2(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_2(\theta^*)\|_2^2 \right] \\ &\quad + 3 \mathbb{E} \left[\|\Sigma^{-1} \nabla^3 F_0(\theta^*) ((\Sigma^{-1} \nabla F_2(\theta^*)) \otimes (\Sigma^{-1} \nabla F_2(\theta^*)))\|_2^2 \right] \\ &\quad + 3r^{-3} \mathcal{O}(n^{-3}). \end{aligned} \quad (51)$$

We now bound the first two terms in inequality (51). Applying the Cauchy-Schwarz inequality and Lemma 2, the first term can be upper bounded as

$$\begin{aligned} &\mathbb{E} \left[\|\Sigma^{-1} (\nabla^2 F_2(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_2(\theta^*)\|_2^2 \right] \\ &\leq \left(\mathbb{E} \left[\|\Sigma^{-1} (\nabla^2 F_2(\theta^*) - \Sigma)\|_2^4 \right] \mathbb{E} \left[\|\Sigma^{-1} \nabla F_2(\theta^*)\|_2^4 \right] \right)^{1/2} \\ &= (r^{-2}) \mathcal{O}(\log^2(d)n^{-2}) \cdot r^{-2} \mathcal{O}(n^{-2})^{1/2} = r^{-2} \mathcal{O}(n^{-2}), \end{aligned}$$

where the order notation subsumes the logarithmic factor in the dimension. Since $\nabla^3 F_0(\theta^*) : \mathbb{R}^{d^2} \rightarrow \mathbb{R}^d$ is linear, the second term in the inequality (51) may be bounded completely analogously as it involves the outer product $\Sigma^{-1} \nabla F_2(\theta^*) \otimes \Sigma^{-1} \nabla F_2(\theta^*)$. (Recall the equality (55) in the proof of Lemma 7.) Recalling the bound (51), we have thus shown that

$$\mathbb{E} \left[\|\theta_2 - \theta^* + \Sigma^{-1} \nabla F_2(\theta^*)\|_2^2 \right] = r^{-2} \mathcal{O}(n^{-2}),$$

or $\theta_2 - \theta^* = -\Sigma^{-1} \nabla F_2(\theta^*) + r^{-1} \mathcal{R}_2$. The proof of the first equality in equation (50) is entirely analogous.

We now apply the equalities (50) to obtain the result of the lemma. We have

$$\mathbb{E} \left[\|\theta_1 - \theta^* - r(\theta_2 - \theta^*)\|_2^2 \right] = \mathbb{E} \left[\|\Sigma^{-1} \nabla F_1(\theta^*) + r \Sigma^{-1} \nabla F_2(\theta^*) + \mathcal{R}_2\|_2^2 \right].$$

Using the inequality $(a + b)^2 \leq (1 + \eta)a^2 + (1 + 1/\eta)b^2$ for any $\eta \geq 0$, we have

$$\begin{aligned} (a + b + c)^2 &\leq (1 + \eta)a^2 + (1 + 1/\eta)(b + c)^2 \\ &\leq (1 + \eta)a^2 + (1 + 1/\eta)(1 + \alpha)b^2 + (1 + 1/\eta)(1 + 1/\alpha)c^2 \end{aligned}$$

for any $\eta, \alpha \geq 0$. Taking $\eta = 1$ and $\alpha = 1/2$, we obtain $(a + b + c)^2 \leq 2a^2 + 3b^2 + 6c^2$, so applying the triangle inequality, we have

$$\begin{aligned} \mathbb{E} \left[\|\theta_1 - \theta^* - r(\theta_2 - \theta^*)\|_2^2 \right] &= \mathbb{E} \left[\left\| -\Sigma^{-1} \nabla F_1(\theta^*) + r \Sigma^{-1} \nabla F_2(\theta^*) + \mathcal{R}_2 \right\|_2^2 \right] \\ &\leq 2\mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^2 \right] + 3r^2 \mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_2(\theta^*) \right\|_2^2 \right] + \mathcal{O}(n^{-2}). \end{aligned} \quad (52)$$

Since F_2 is a sub-sampled version of F_1 , algebraic manipulations yield

$$\mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_2(\theta^*) \right\|_2^2 \right] = \frac{n}{rn} \mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^2 \right] = \frac{1}{r} \mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^2 \right]. \quad (53)$$

Combining equations (52) and (53), we obtain the desired bound (37).

H Proof of Lemma 7

The proof follows from a slightly more careful application of the Taylor expansion (23). The starting point in our proof is to recall the success events (22) and the joint event $\mathcal{E} := \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$. We begin by arguing that we may focus on the case where \mathcal{E} holds. Let C denote the right hand side of the equality (38) except for the remainder \mathcal{R}_3 term. By Assumption C, we follow the bound (27) to find that

$$\mathbb{E} \left[1_{(\mathcal{E}^c)} \|\theta_1 - \theta^*\|_2^2 \right] = \mathcal{O}(R^2 n^{-4}),$$

so we can focus on the case where the joint event $\mathcal{E} = \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$ does occur.

Defining $\Delta = \theta_1 - \theta^*$ for notational convenience, on \mathcal{E} we have that for some $\kappa \in [0, 1]$, with $\theta' = (1 - \kappa)\theta_1 + \kappa\theta^*$,

$$\begin{aligned} 0 &= \nabla F_1(\theta^*) + \nabla^2 F_1(\theta^*) \Delta + \nabla^3 F_1(\theta') (\Delta \otimes \Delta) \\ &= \nabla F_1(\theta^*) + \nabla^2 F_0(\theta^*) \Delta + \nabla^3 F_0(\theta^*) (\Delta \otimes \Delta) \\ &\quad + (\nabla^2 F_1(\theta^*) - \nabla^2 F_0(\theta^*)) \Delta + (\nabla^3 F_1(\theta') - \nabla^3 F_0(\theta^*)) (\Delta \otimes \Delta). \end{aligned}$$

Now, we recall the definition $\Sigma = \nabla^2 F_0(\theta^*)$, the Hessian of the risk at the optimal point, and solve for the error Δ to see that

$$\begin{aligned} \Delta &= -\Sigma^{-1} \nabla F_1(\theta^*) - \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Delta - \Sigma^{-1} \nabla^3 F_1(\theta') (\Delta \otimes \Delta) \\ &\quad + \Sigma^{-1} (\nabla^3 F_0(\theta^*) - \nabla^3 F_1(\theta')) (\Delta \otimes \Delta) \end{aligned} \quad (54)$$

on the event \mathcal{E} . As we did in the proof of Theorem 1, specifically in deriving the recursive equality (29), we may apply the expansion (25) of $\Delta = \theta_1 - \theta^*$ to obtain a clean asymptotic expansion of Δ using (54). Recall the definition $P = \nabla^2 F_0(\theta^*) - \nabla^2 F_1(\theta^*)$ for shorthand here (as in the expansion (25), though we no longer require Q).

First, we claim that

$$1_{(\mathcal{E})} (\nabla^3 F_0(\theta^*) - \nabla^3 F_1(\theta')) (\Delta \otimes \Delta) = \mathcal{R}_3. \quad (55)$$

Adding and subtracting $\nabla^3 F_1(\theta^*)$ from the above expression (and dropping $1_{(\mathcal{E})}$ for simplicity) we must control

$$(\nabla^3 F_0(\theta^*) - \nabla^3 F_1(\theta^*)) (\Delta \otimes \Delta) + (\nabla^3 F_1(\theta^*) - \nabla^3 F_1(\theta')) (\Delta \otimes \Delta).$$

To begin, recall that $\|u \otimes v\|_2 = \|uv^\top\|_2 = \|u\|_2 \|v\|_2$. By Assumption D, on the event \mathcal{E} we have that $\nabla^3 F_1$ is $(1/n) \sum_{i=1}^n M(X_i)$ -Lipschitz, so defining $M_n = (1/n) \sum_{i=1}^n M(X_i)$, we have

$$\begin{aligned} \mathbb{E} \left[1_{(\mathcal{E})} \|(\nabla^3 F_1(\theta^*) - \nabla^3 F_1(\theta'))(\Delta \otimes \Delta)\|_2^2 \right] &\leq \mathbb{E} \left[M_n \|\theta^* - \theta'\|_2 \|\Delta \otimes \Delta\|_2^2 \right] \\ &\leq \mathbb{E} [M_n^4]^{1/4} \mathbb{E} [\|\theta_1 - \theta^*\|_2^8]^{3/4} = \mathcal{O}(n^{-3}) \end{aligned}$$

by Hölder's inequality and Lemma 3. The remaining term we must control is the derivative difference $\mathbb{E}[\|(\nabla^3 F_1(\theta^*) - \nabla^3 F_0(\theta^*))(\Delta \otimes \Delta)\|_2^2]$. Define the random vector-valued function $\mathbf{G} = \nabla(F_1 - F_0)$, and let \mathbf{G}_j denote its j th coordinate. Then by definition we have

$$(\nabla^3 F_1(\theta^*) - \nabla^3 F_0(\theta^*))(\Delta \otimes \Delta) = \left[\Delta^\top (\nabla^2 \mathbf{G}_1(\theta^*)) \Delta \quad \cdots \quad \Delta^\top (\nabla^2 \mathbf{G}_d(\theta^*)) \Delta \right]^\top \in \mathbb{R}^d.$$

Therefore, by the Cauchy-Schwarz inequality and the fact that $x^\top A x \leq \|A\|_2 \|x\|_2^2$,

$$\begin{aligned} \mathbb{E} \left[\|(\nabla^3 F_1(\theta^*) - \nabla^3 F_0(\theta^*))(\Delta \otimes \Delta)\|_2^2 \right] &= \sum_{j=1}^d \mathbb{E} \left[\left(\Delta^\top (\nabla^2 \mathbf{G}_j(\theta^*)) \Delta \right)^2 \right] \\ &\leq \sum_{j=1}^d \left(\mathbb{E} [\|\Delta\|_2^8] \mathbb{E} [\|\nabla^2 \mathbf{G}_j(\theta^*)\|_2^4] \right)^{1/2}. \end{aligned}$$

Applying Lemma 3 yields that $\mathbb{E}[\|\Delta\|_2^8] = \mathcal{O}(n^{-4})$. By defining the functions $\mathbf{g}(\cdot; x) := \nabla f(\cdot; x) - \nabla F_0(\cdot)$, we can write

$$\nabla^2 \mathbf{G}_j(\theta^*) = \frac{1}{n} \sum_{i=1}^n \nabla^2 \mathbf{g}_j(\theta^*; X_i)$$

For every coordinate j , the random matrices $\nabla^2 \mathbf{g}_j(\theta^*; X_i)$ ($i = 1, \dots, n$) are i.i.d. and mean zero. By Assumption C, we have $\|\nabla^2 \mathbf{g}_j(\theta^*; X_i)\|_2 \leq 2L(X_i)$, whence $\mathbb{E}[\|\nabla^2 \mathbf{g}_j(\theta^*; X_i)\|_2^8] \leq 2^8 L^8$. Applying Lemma 10 (or the technique used to prove inequality (26b) following Lemma 10), we thus obtain

$$\mathbb{E} \left[\|\nabla^2 \mathbf{G}_j(\theta^*)\|_2^4 \right] = \mathcal{O}(n^{-2} \log^2(d)),$$

so

$$\mathbb{E} \left[\|(\nabla^3 F_1(\theta^*) - \nabla^3 F_0(\theta^*))(\Delta \otimes \Delta)\|_2^2 \right] = \mathcal{O}(d \log(d) n^{-3}).$$

In particular, we have our desired result (55).

Now we claim that

$$1_{(\mathcal{E})} \nabla^3 F_1(\theta^*)(\Delta \otimes \Delta) = \nabla^3 F_1(\theta^*)((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) + \mathcal{R}_3. \quad (56)$$

Indeed, applying the expansion (25) to the difference $\Delta = \theta_1 - \theta^*$, we have on \mathcal{E} that

$$\begin{aligned} \Delta \otimes \Delta &= (\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*)) + (\Sigma^{-1} P \Delta) \otimes (\Sigma^{-1} P \Delta) \\ &\quad - (\Sigma^{-1} P \Delta) \otimes (\Sigma^{-1} \nabla F_1(\theta^*)) - (\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} P \Delta). \end{aligned}$$

We can bound each of the second three outer products in the equality above similarly; we focus on the last for simplicity. Applying the Cauchy-Schwarz inequality, we have

$$\mathbb{E} \left[\left\| (\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} P \Delta) \right\|_2^2 \right] \leq \left(\mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^4 \right] \mathbb{E} \left[\left\| \Sigma^{-1} P(\theta_1 - \theta^*) \right\|_2^4 \right] \right)^{\frac{1}{2}}.$$

From Lemmas 3 and 4, we obtain that

$$\mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^4 \right] = \mathcal{O}(n^{-2}) \quad \text{and} \quad \mathbb{E} \left[\left\| \Sigma^{-1} P(\theta_1 - \theta^*) \right\|_2^4 \right] = \mathcal{O}(n^{-4})$$

after an additional application of Cauchy-Schwarz for the second expectation. This shows that

$$(\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} P \Delta) = \mathcal{R}_3,$$

and a similar proof applies to the other three terms in the outer product $\Delta \otimes \Delta$. Using the linearity of $\nabla^3 F_1(\theta^*)$, we see that to prove the equality (56), all that is required is that

$$1_{(\mathcal{E}^c)} \nabla^3 F_1(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) = \mathcal{R}_3. \quad (57)$$

For this, we apply Hölder's inequality several times. Indeed, we have

$$\begin{aligned} & \mathbb{E} \left[\left\| 1_{(\mathcal{E}^c)} \nabla^3 F_1(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) \right\|_2^2 \right] \\ & \leq \mathbb{E}[1_{(\mathcal{E}^c)}]^{1/4} \mathbb{E} \left[\left\| \nabla^3 F_1(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) \right\|_2^{8/3} \right]^{3/4} \\ & \leq \mathbb{E}[1_{(\mathcal{E}^c)}]^{1/4} \mathbb{E} \left[\left\| \nabla^3 F_1(\theta^*) \right\|^{8/3} \left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^{16/3} \right]^{3/4} \\ & \leq \mathbb{E}[1_{(\mathcal{E}^c)}]^{1/4} \mathbb{E} \left[\left\| \nabla^3 F_1(\theta^*) \right\|^8 \right]^{1/4} \mathbb{E} \left[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^8 \right]^{2/4} = \mathcal{O}(n^{-1} \cdot L^2 \cdot n^{-2}). \end{aligned}$$

For the final asymptotic bound, we used Eq. (27) to bound $\mathbb{E}[1_{(\mathcal{E}^c)}]$, used the fact (from Assumption C) that $\mathbb{E}[L(X)^8] \leq L^8$ to bound the term involving $\nabla^3 F_1(\theta^*)$, and applied Lemma 3 to control $\mathbb{E}[\left\| \Sigma^{-1} \nabla F_1(\theta^*) \right\|_2^8]$. Thus the equality (57) holds, and this completes the proof of the equality (56).

For the final step in the lemma, we claim that

$$-1_{(\mathcal{E})} \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Delta = \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*) + \mathcal{R}_3. \quad (58)$$

To prove (58) requires an argument completely parallel to that for our claim (56). As before, we use the expansion (25) of the difference Δ to obtain that on \mathcal{E} ,

$$\begin{aligned} & -\Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Delta \\ & = \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*) - \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} P \Delta. \end{aligned}$$

Now apply Lemmas 3 and 4 to the final term after a few applications of Hölder's inequality. To finish the equality (58), we argue that $1_{(\mathcal{E}^c)} \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*) = \mathcal{R}_3$, which follows exactly the line of reasoning used to prove the remainder (57).

Applying equalities (55), (56), and (58) to our earlier expansion (54) yields that

$$\begin{aligned} \Delta & = 1_{(\mathcal{E})} \left[-\Sigma^{-1} \nabla F_1(\theta^*) - \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Delta - \Sigma^{-1} \nabla^3 F_1(\theta^*) (\Delta \otimes \Delta) \right. \\ & \quad \left. + \Sigma^{-1} (\nabla^3 F_0(\theta^*) - \nabla^3 F_1(\theta')) (\Delta \otimes \Delta) \right] + 1_{(\mathcal{E}^c)} \Delta \\ & = -\Sigma^{-1} \nabla F_1(\theta^*) + \Sigma^{-1} (\nabla^2 F_1(\theta^*) - \Sigma) \Sigma^{-1} \nabla F_1(\theta^*) \\ & \quad - \Sigma^{-1} \nabla^3 F_1(\theta^*) ((\Sigma^{-1} \nabla F_1(\theta^*)) \otimes (\Sigma^{-1} \nabla F_1(\theta^*))) + \mathcal{R}_3 + 1_{(\mathcal{E}^c)} \Delta. \end{aligned}$$

Of course, $\mathbb{E}[1_{(\mathcal{E}^c)} \left\| \Delta \right\|_2^2] \leq \mathbb{P}(\mathcal{E}^c) R^2 = \mathcal{O}(n^{-4})$, so this yields the desired statement of the lemma. \square

References

- [1] A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 25*, 2011.
- [2] A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. J. Wainwright. Information-theoretic lower bounds on the oracle complexity of convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, May 2012.
- [3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] R. Chen, A. Gittens, and J. A. Tropp. The masked sample covariance estimator: an analysis using matrix concentration inequalities. *Information and Inference*, to appear, 2012.
- [5] A. de Acosta. Inequalities for B -valued random vectors with applications to the strong law of large numbers. *The Annals of Probability*, 9:157–161, 1981.
- [6] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13:165–202, 2012.
- [7] J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization: convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2012.
- [8] J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- [9] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall, 1993.
- [10] P. Hall. *The Bootstrap and Edgeworth Expansion*. Springer, 1992.
- [11] E. Hazan, A. Kalai, S. Kale, and A. Agarwal. Logarithmic regret algorithms for online convex optimization. In *Proceedings of the Nineteenth Annual Conference on Computational Learning Theory*, 2006.
- [12] B. Johansson, M. Rabi, and M. Johansson. A randomized incremental subgradient method for distributed optimization in networked systems. *SIAM Journal on Optimization*, 20(3):1157–1170, 2009.
- [13] R. W. Keener. *Theoretical Statistics: Topics for a Core Course*. Springer, 2010.
- [14] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Springer, 1991.
- [15] E. L. Lehmann and G. Casella. *Theory of Point Estimation, Second Edition*. Springer, 1998.
- [16] G. Mann, R. McDonald, M. Mohri, N. Silberman, and D. Walker. Efficient Large-Scale Distributed Training of Conditional Maximum Entropy Models. In *Advances in Neural Information Processing Systems 22*, pages 1231–1239, 2009.
- [17] C. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.

- [18] R. McDonald, K. Hall, and G. Mann. Distributed training strategies for the structured perceptron. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2010.
- [19] A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54:48–61, 2009.
- [20] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [21] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.
- [22] D. N. Politis, J. P. Romano, and M. Wolf. *Subsampling*. Springer, 1999.
- [23] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [24] S. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [25] S. S. Ram, A. Nedić, and V. V. Veeravalli. Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of Optimization Theory and Applications*, 147(3): 516–545, 2010.
- [26] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: a lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems 25*, 2011.
- [27] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- [28] G. Sun. KDD cup track 2 soso.com ads prediction challenge, 2012. URL <http://www.kddcup2012.org/c/kddcup2012-track2>. Accessed August 1, 2012.
- [29] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998. ISBN 0-521-49603-9.
- [30] M. A. Zinkevich, A. Smola, M. Weimer, and L. Li. Parallelized Stochastic Gradient Descent. In *Advances in Neural Information Processing Systems 24*, 2010.